A PennWell Publication

AUGUST 1984

COMPUTER DESIGN

THE MAGAZINE OF COMPUTER BASED SYSTEMS

SEMICONDUCTOR MEMORIES



A PennWell Publication

AUGUST 1984

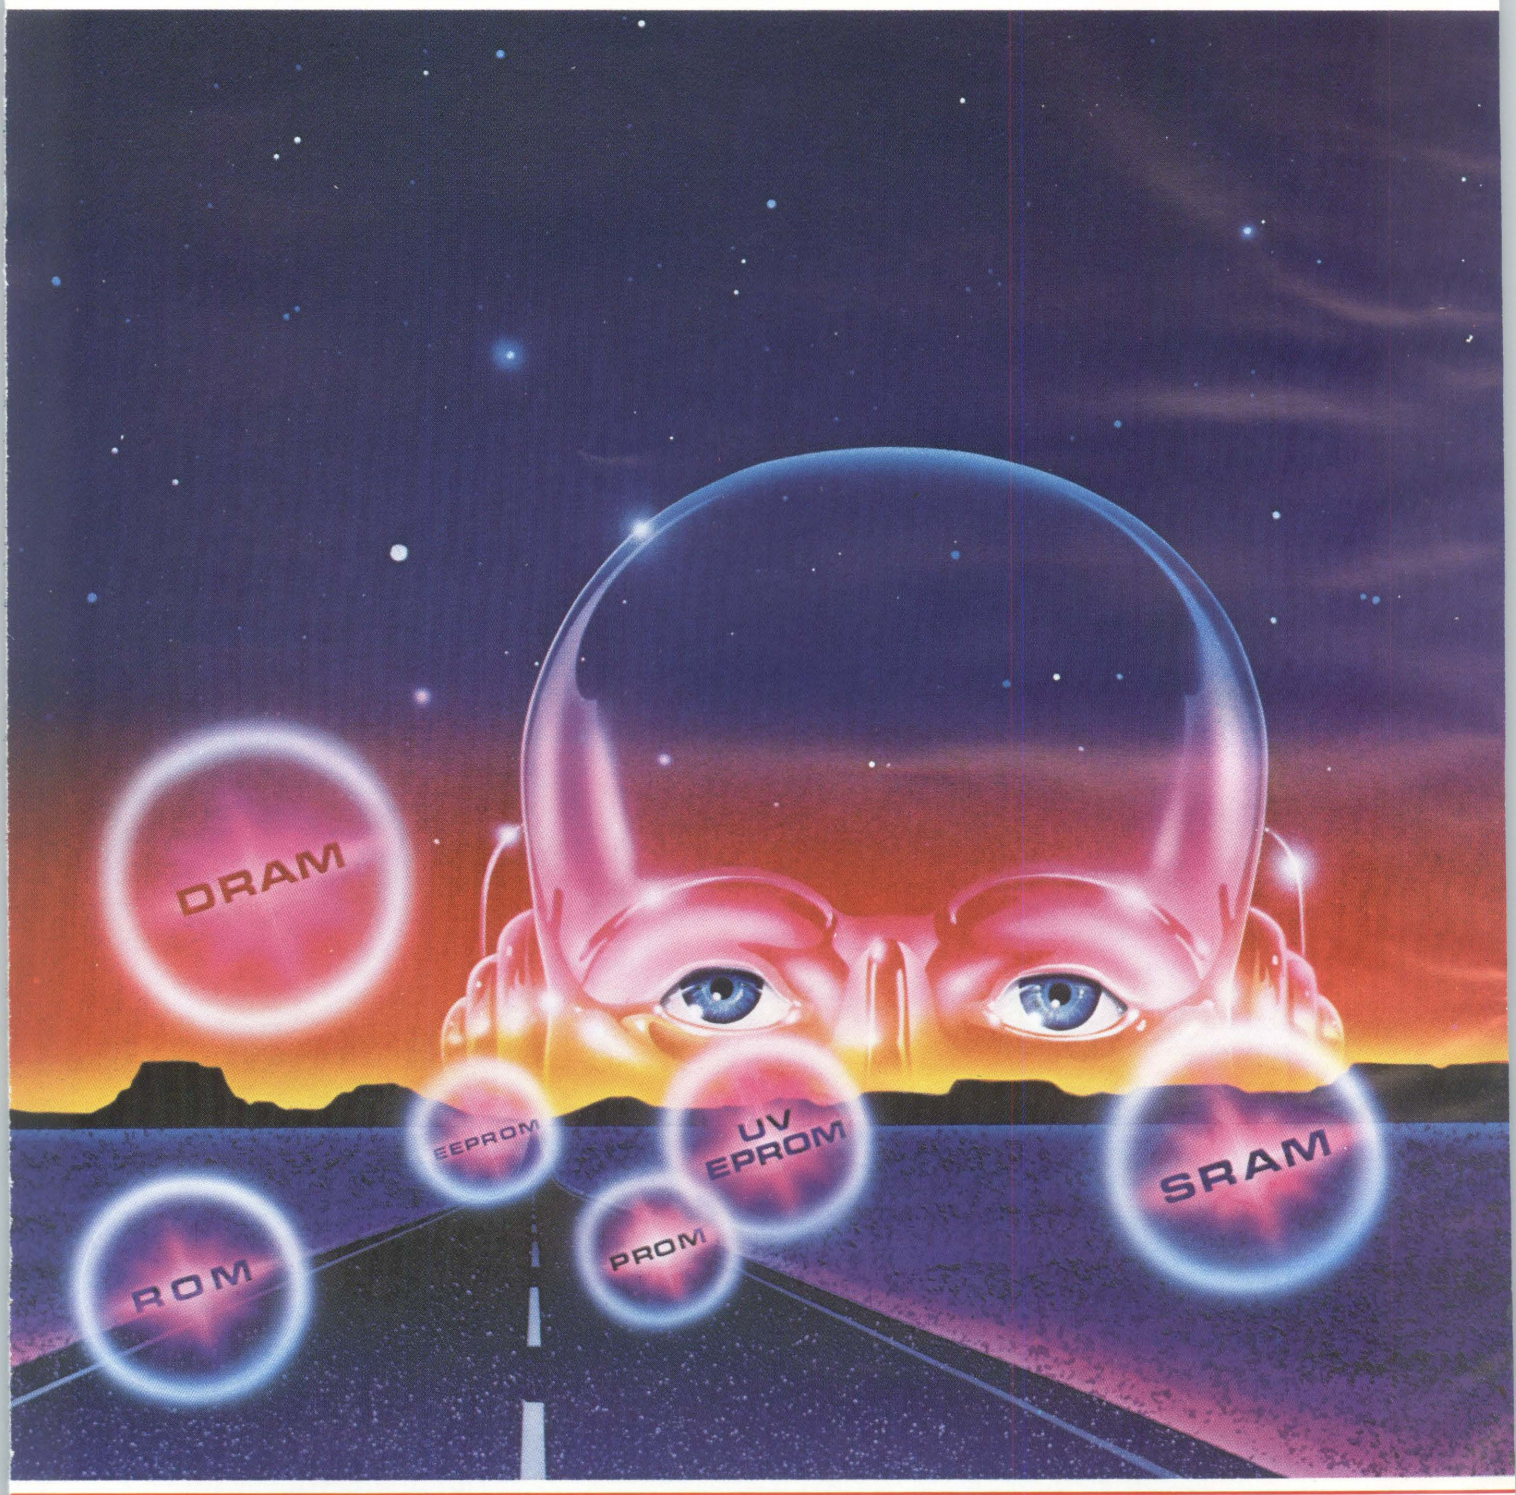
COMPUTER DESIGN

THE MAGAZINE OF COMPUTER BASED SYSTEMS

SEMICONDUCTOR MEMORIES



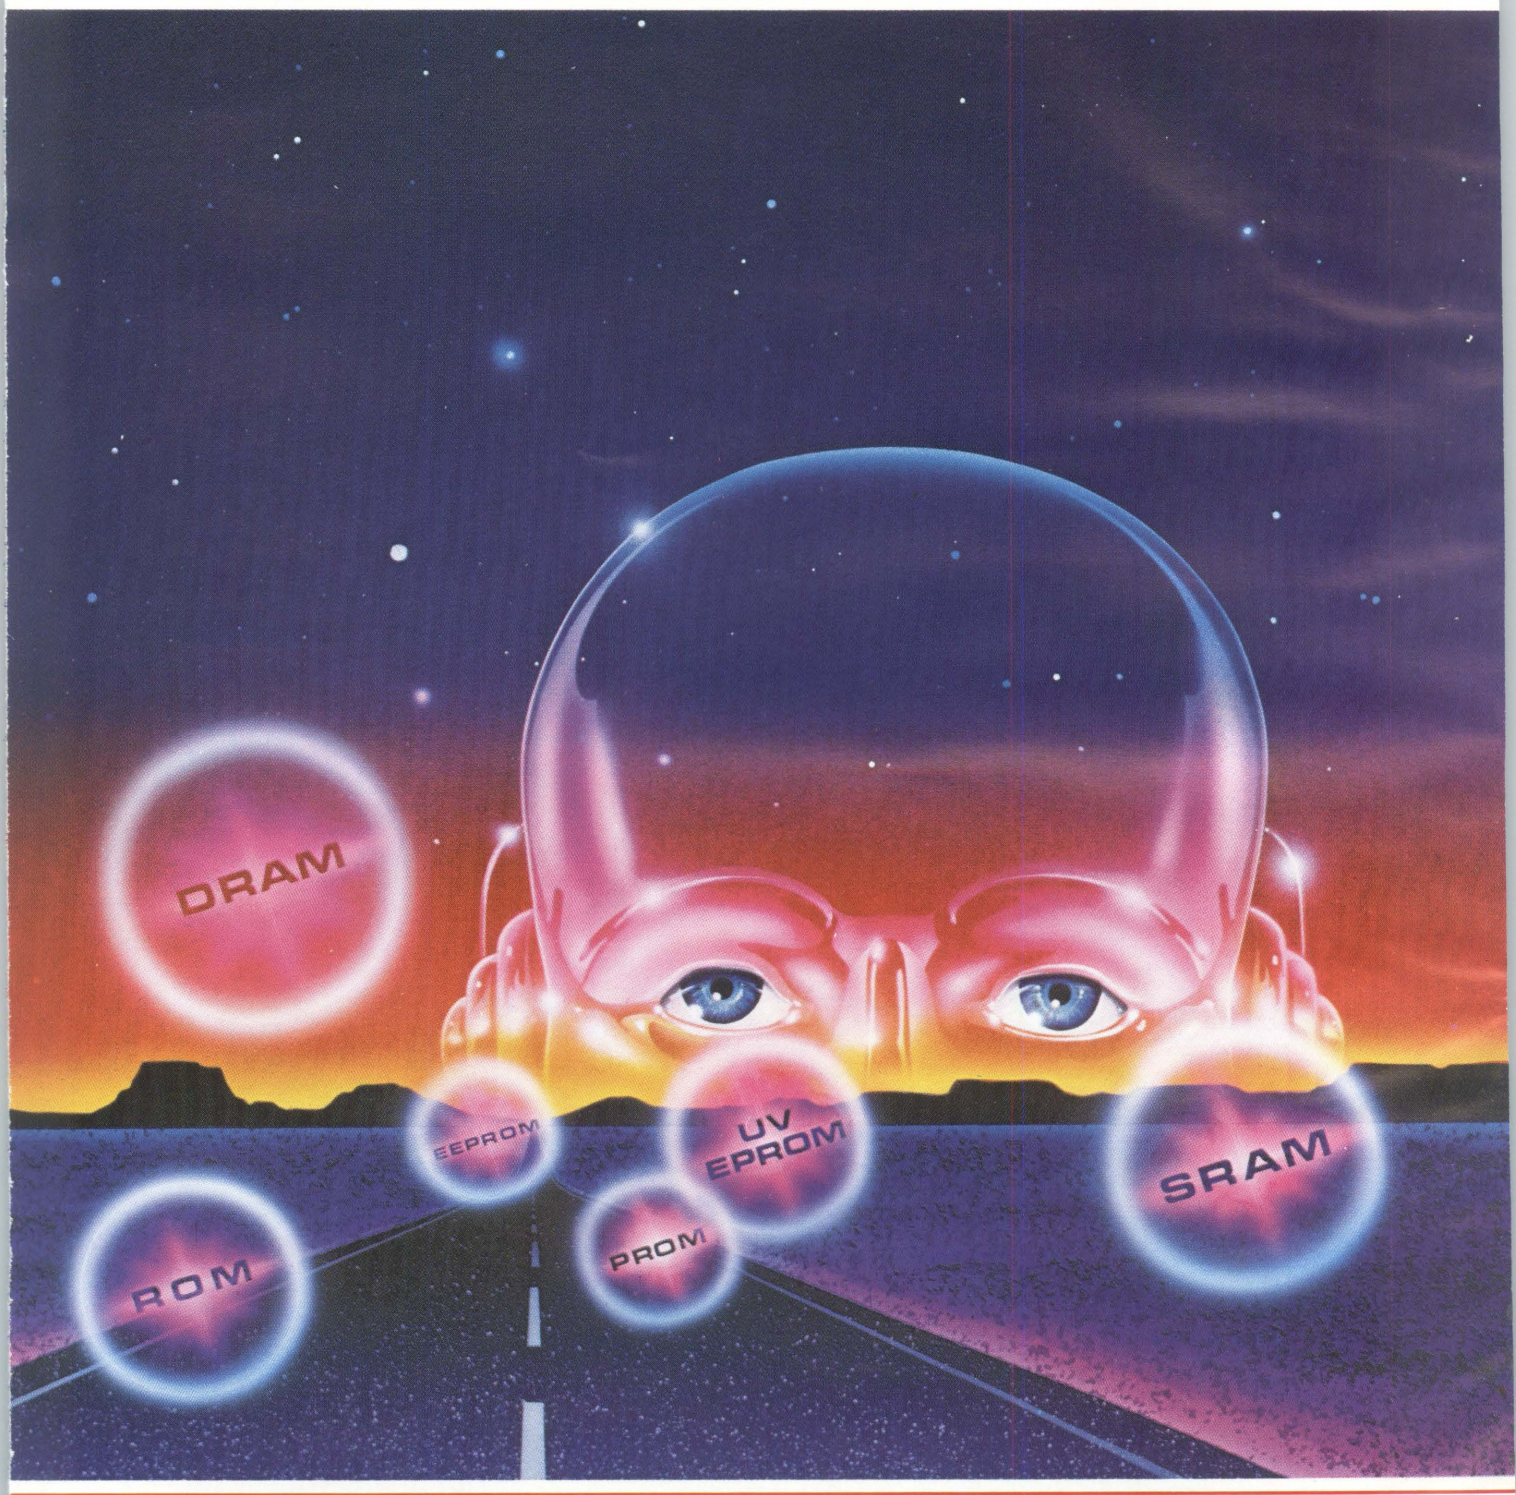
IN OUR MEMORIES.



CIRCLE 1

NEC
NEC Electronics

IN OUR MEMORIES.



CIRCLE 1

NEC
NEC Electronics

THERE'S NO LAPSE

AT NEC, YOU'LL FIND QUALITY MEMORIES IN EVERY CATEGORY.

All in one place.

Chances are, we have the exact memories you're looking for. Leading-edge 25ns 4K static RAMs. 64K static RAMs. 128K UV EPROMs. New 256K DRAMs. And our remarkable one *megabit* ROM—the highest density ROM in the world.

These aren't promises of things to come. They're here. Now. And ready for your next design.

You can trust our memories.

At NEC, quality is understood. Fact is, you can expect 100% burn-in, standard. And a guaranteed AQL of 0.1%. Which means, with NEC, your memories will be worry-free.

A complete package deal.

Be reminded, our memory devices come in a variety of packages. Standard and skinny. Plastic and ceramic DIPs. Flat packages and leadless chip carriers.

So if you have designs on the future, remember NEC. For a never-ending line of quality memories.

NEC Memory Products

Type of Memory	Technologies				Densities								
	CMOS	MMOS[1]	NMOS	B/P	1K	2K	4K	8K	16K	32K	64K	128K	256K	1M
Static RAM[2]	•	•	•				•		•		•			
Dynamic RAM		•	•				•		•		•		•	
ROM	•		•							•	•	•	•	•
PROM				•	•		•	•	•					
EEPROM			•						•					
UV EPROM	•		•						•	•	•	•		

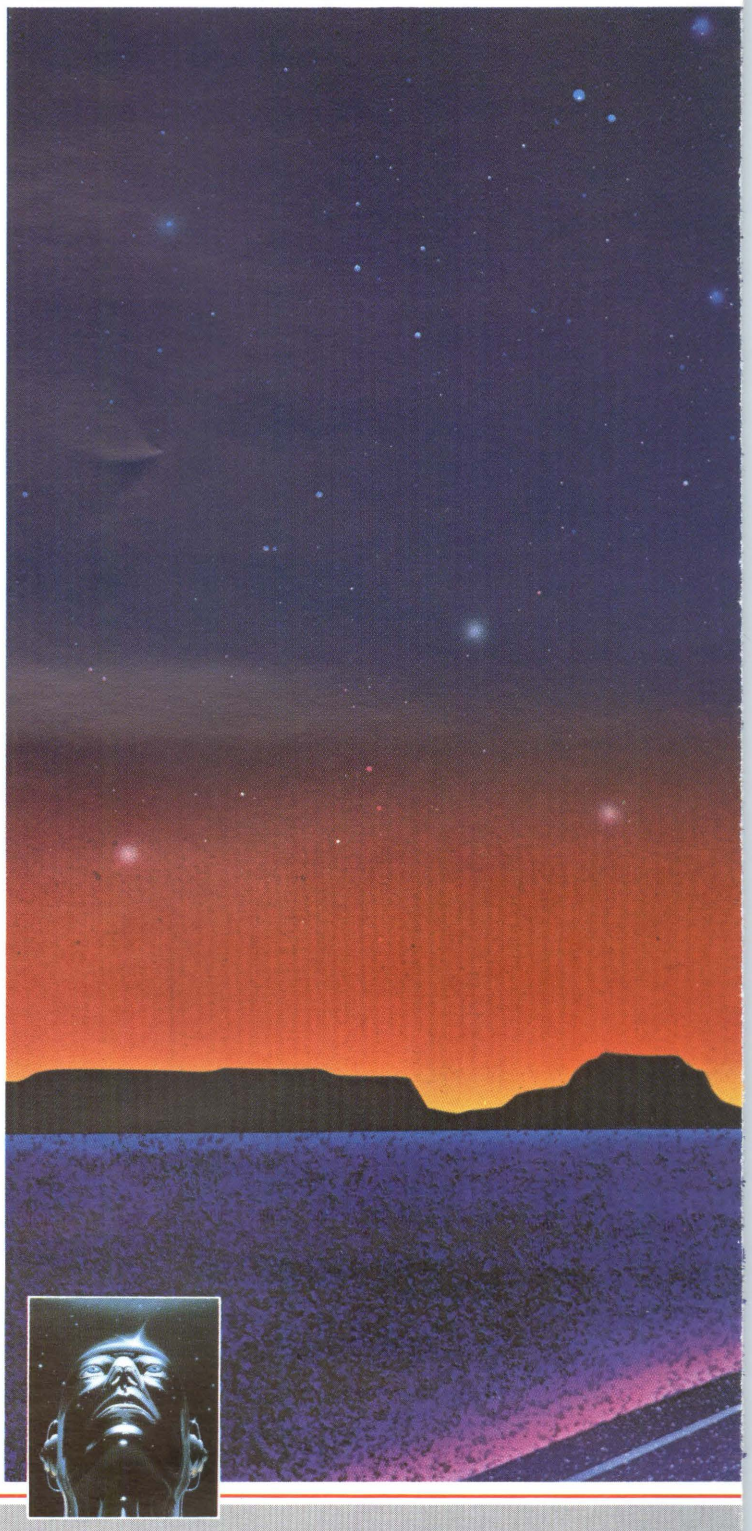
[1]MMOS refers to Mixed-MOS technology (CMOS and NMOS).
[2]High speed 4K and 16K devices available.

WE'RE TAKING ON THE FUTURE.

For a fast response about NEC Electronics' complete line of memory products, call TOLL FREE 1-800-556-1234, ext. 188. In California, call 1-800-441-2345, ext. 188.

NEC sales offices: Woburn, MA (617) 935-6339 · Melville, NY (516) 423-2500 · Poughkeepsie, NY (914) 452-4747 · Pompano Beach, FL (305) 785-8250 · Columbia, MD (301) 730-8600 · Norcross, GA (404) 447-4409 · Arlington Heights, IL (312) 577-9090 · Southfield, MI (313) 559-4242 · Bloomington, MN (612) 854-4443 · Dallas, TX (214) 931-0641 · Orange, CA (714) 937-5244 · Cupertino, CA (408) 446-0650

© 1983, NEC Electronics



THERE'S NO LAPSE

AT NEC, YOU'LL FIND QUALITY MEMORIES IN EVERY CATEGORY.

All in one place.

Chances are, we have the exact memories you're looking for. Leading-edge 25ns 4K static RAMs. 64K static RAMs. 128K UV EPROMs. New 256K DRAMs. And our remarkable one *megabit* ROM—the highest density ROM in the world.

These aren't promises of things to come. They're here. Now. And ready for your next design.

You can trust our memories.

At NEC, quality is understood. Fact is, you can expect 100% burn-in, standard. And a guaranteed AQL of 0.1%. Which means, with NEC, your memories will be worry-free.

A complete package deal.

Be reminded, our memory devices come in a variety of packages. Standard and skinny. Plastic and ceramic DIPs. Flat packages and leadless chip carriers.

So if you have designs on the future, remember NEC. For a never-ending line of quality memories.

NEC Memory Products

Type of Memory	Technologies				Densities								
	CMOS	MMOS[1]	NMOS	B/P	1K	2K	4K	8K	16K	32K	64K	128K	256K	1M
Static RAM[2]	•	•	•				•		•		•			
Dynamic RAM		•	•				•		•		•		•	
ROM	•		•							•	•	•	•	•
PROM				•	•		•	•	•					
EEPROM			•						•					
UV EPROM	•		•						•	•	•	•		

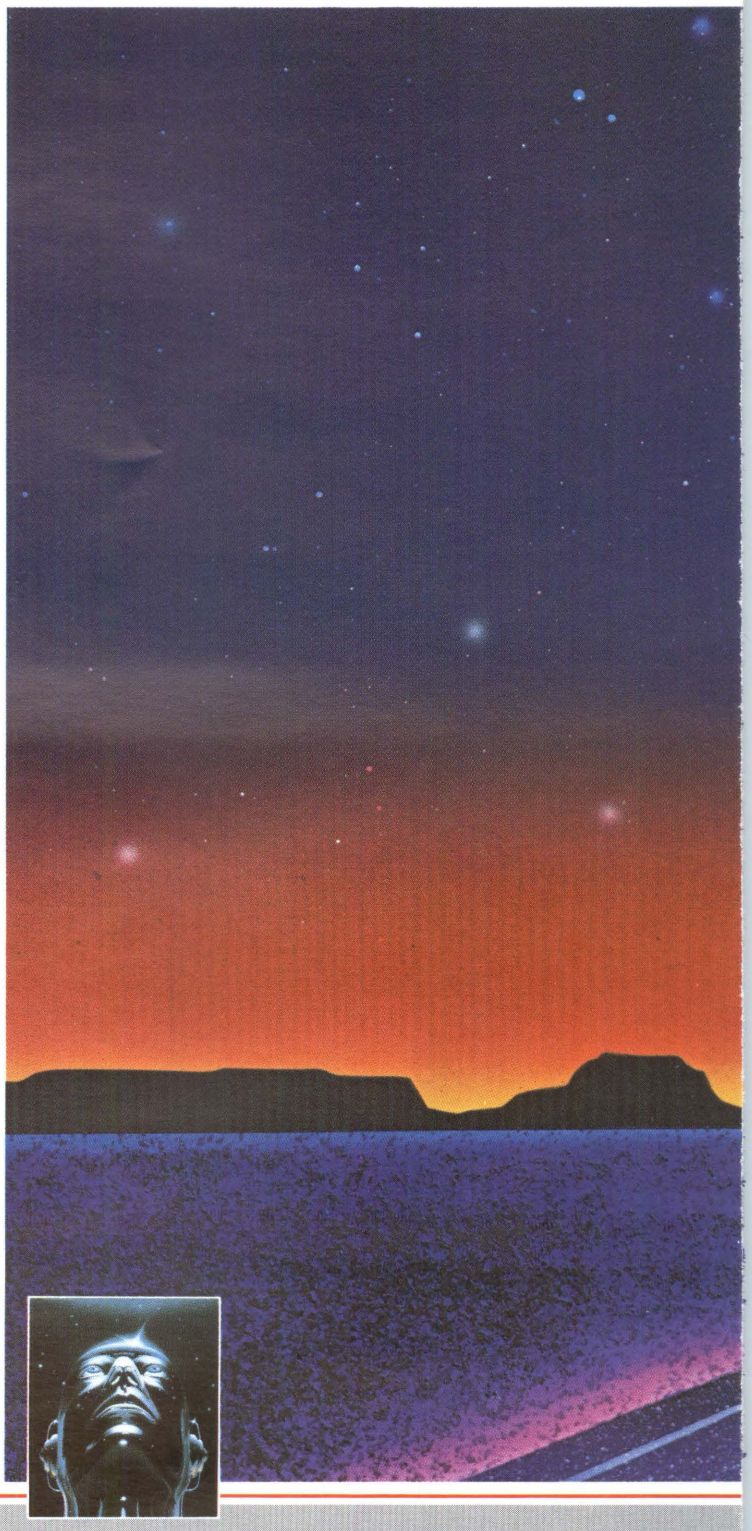
[1]MMOS refers to Mixed-MOS technology (CMOS and NMOS).
[2]High speed 4K and 16K devices available.

WE'RE TAKING ON THE FUTURE.

For a fast response about NEC Electronics' complete line of memory products, call TOLL FREE 1-800-556-1234, ext. 188. In California, call 1-800-441-2345, ext. 188.

NEC sales offices: Woburn, MA (617) 935-6339 · Melville, NY (516) 423-2500 · Poughkeepsie, NY (914) 452-4747 · Pompano Beach, FL (305) 785-8250 · Columbia, MD (301) 730-8600 · Norcross, GA (404) 447-4409 · Arlington Heights, IL (312) 577-9090 · Southfield, MI (313) 559-4242 · Bloomington, MN (612) 854-4443 · Dallas, TX (214) 931-0641 · Orange, CA (714) 937-5244 · Cupertino, CA (408) 446-0650

© 1983, NEC Electronics



MEMORIES OF THE FUTURE



MEMORIES OF THE FUTURE



Model 6470 Cartridge Recorder — the latest distinguished member of a

PIONEER FAMILY

Kennedy didn't just pioneer the ¼" cartridge field — it started it, with the first practical recorder using the new ¼" 3M cartridge.

Model 6470 continues the tradition with the usual package of new features, such as:

- High Capacity — Model 6470 can store up to 54 MBYTES of formatted data on one cartridge.
- Start/Stop with data streamings at 37.5 ips.
- Full Command Repertoire, utilizing either the Pico Bus or Pertec interface.
- Automatic Read Thresholds — during read retry three different thresholds are automatically selected, a feature usually found only on large 9 track tape transports.
- Backward Compatibility — Model 6470 will read tapes written by Model 6455.

These are but a few of Model 6470's many features. Write or give us a call today for the complete story of the newest member of this oldest family.

KENNEDY
An Allegheny International Company
1600 Shamrock Ave., Monrovia, CA 91016
(818) 357-8831 • ITT TELEX 472-0116 KENNEDY



KENNEDY • QUALITY • COUNT ON IT

CIRCLE 2

Model 6470 Cartridge Recorder — the latest distinguished member of a

PIONEER FAMILY

Kennedy didn't just pioneer the ¼" cartridge field — it started it, with the first practical recorder using the new ¼" 3M cartridge.

Model 6470 continues the tradition with the usual package of new features, such as:

- High Capacity — Model 6470 can store up to 54 MBYTES of formatted data on one cartridge.
- Start/Stop with data streamings at 37.5 ips.
- Full Command Repertoire, utilizing either the Pico Bus or Pertec interface.
- Automatic Read Thresholds — during read retry three different thresholds are automatically selected, a feature usually found only on large 9 track tape transports.
- Backward Compatibility — Model 6470 will read tapes written by Model 6455.

These are but a few of Model 6470's many features. Write or give us a call today for the complete story of the newest member of this oldest family.

KENNEDY
An Allegheny International Company
1600 Shamrock Ave., Monrovia, CA 91016
(818) 357-8831 • ITT TELEX 472-0116 KENNEDY

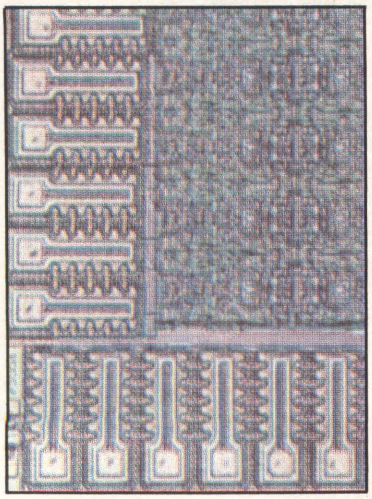


KENNEDY • QUALITY • COUNT ON IT

CIRCLE 2

COMPUTER DESIGN®

System technology



Page 29

29 Packaging & power:
Power and logic devices are merging on the same chip

38 Computers:
Matrix array processor breaks through supercomputer barrier

44 Software:
Big Blue makes a major move to link computers

53 Software:
Multitasking Unix look-alike supports 11 IBM XT users

60 Peripherals:
Flat panels approach monochrome CRT specifications

System design

77 Packaging & power: **Connectors—the missing link in EMI suppression**
by Frank Drzymkowski and Dave Goodman—Computer equipment must meet stringent FCC restrictions on rf emissions. Although the first line of defense for designers is proper packaging, electromagnetic interference can occur between units if attention is not paid to connector design.

91 Software: **Software quality: design it in from the start**
by M. Ghiassi—A comprehensive software quality assurance program should be an integral part of the development process. This program ensures not only that functions work as planned, but that planned functions are useful.

Midcon/84 and Mini/Micro Southwest

65 Midcon and Mini/Micro Southwest, two major southwest OEM electronics and computer conferences, are being held concurrently this year in Dallas, Texas on September 11-13. Keynote speaker Admiral "Bobby" Inman, CED of the Microelectronics and Computer Technology Group, will lead off the two programs of technical sessions. Of specific interest to computer system designers are the technical sessions at Mini/Micro as well as many at Midcon.

Vol 23, No. 9
Aug 1984

Special report on semiconductor memories

101 Fed by breakthroughs in CMOS and the insatiable demand of modern systems for memory, semiconductor memory technology is branching beyond mere "by 1" density increases. VLSI memories are providing system solutions and breaking down performance bottlenecks. In addition to more efficient, denser main memories, RAMs, EPROMs, and EEPROMs are appearing in the parts of the system where their special features are most needed and offering new design possibilities.

Special report on minicomputer operating systems

161 Minicomputer system designers will find a Unix operating system or a Unix look-alike often meets their needs for increased operating system functionality. In fact, the latest minicomputer operating systems for the past five years are more like operating environments and sport both realtime capabilities and fault-tolerant features—whether or not they are Unix based. Computer system designers opting for minicomputers need to know how minicomputer operating systems determine their computer's behavior.



This month's cover was prepared by Mark Lindquist on the Digital Effects Video Palette III and D-48 high resolution camera system. The data base was constructed by Joe Pasquale. The cover was based upon material supplied by United Technologies Mostek.

System components

233 Test & development:
Workstation graphics architecture lightens programming load

234 Integrated circuits:
Dynamic RAMs deliver 256-Kbit power using CMOS process

234 Computers:
High end supermini melds Unix with marked speed

236 Software:
Structured analysis eases project specification

236 Computers:
Unix systems provide long-term solutions with upgradeability

237 Memory systems:
Winchester merges with controller for single-board integration

237 Integrated circuits:
Gate array series reaches 11,000-circuit density

Departments

5 Up front
13 Editorial
23 Letters to the editor
273 Literature
274 Calendar
277 Designer's bookcase
279 System showcase
280 Advertisers' index
283 Reader inquiry card
283 Change of address card

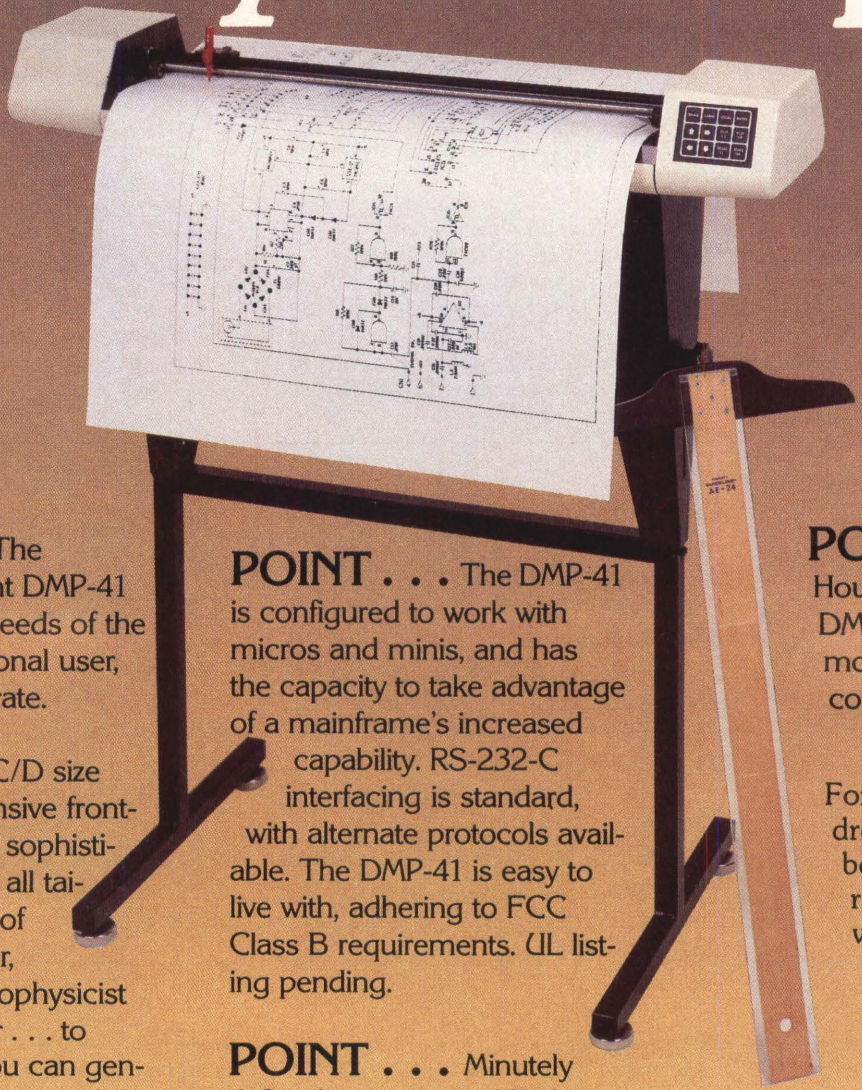
COMPUTER DESIGN©1984 (ISSN-0010-4566) is published monthly, with a thirteenth and fourteenth issue respectively in June and October by PennWell Publishing Company, Advanced Technology Group, 119 Russell Street, Littleton, MA 01460. Second-class postage paid at Littleton, MA 01460 and additional mailing offices. COMPUTER DESIGN is distributed without charge to U.S. and W. Europe-based engineers and engineering managers responsible for computer-based equipment and systems design. Subscription rate for others is $50 in U.S.A. and $75 elsewhere. Single copy price is $5.00 in U.S.A. and $7.50 elsewhere. Microfilm copies of COMPUTER DESIGN are available and may be purchased from University Microfilms, a Xerox Company, 300 North Zeeb Road, Ann Arbor, Michigan 48106. POSTMASTER: CHANGE OF ADDRESS-FORM 3579 to be sent to COMPUTER DESIGN, Circulation Department, P.O. Box 593, Littleton, MA 01460 (USPS 127-340).

Officers of PennWell Publishing Company, 1421 S. Sheridan, Tulsa, OK 74101: P. C. Lauinger, Chairman; Philip C. Lauinger, Jr., President; Joseph A. Wolking, Senior Vice President; H. Mason Fackert, Group Vice President; Carl J. Lawrence, Group Vice President; V. John Maney, Vice President/Finance; L. John Ford, Vice President.

VBPA ☆ABP

® COMPUTER DESIGN is a registered trademark. All rights reserved. No materials may be reprinted without permission. Phone (617) 486-9501.

It's easy to make points when you're a pro.



POINT . . . The Houston Instrument DMP-41 plotter meets the needs of the serious or professional user, yet it's easy to operate.

POINT . . . C/D size format, comprehensive front-panel controls and sophisticated firmware are all tailored to the needs of the surveyor, drafter, oceanographer, geophysicist and land developer . . . to name but a few. You can generate superior architectural elevations, contour maps, circuit-board layouts and assembly drawings quickly and accurately on bond, vellum or synthetic media.

POINT . . . The DMP-41 is configured to work with micros and minis, and has the capacity to take advantage of a mainframe's increased capability. RS-232-C interfacing is standard, with alternate protocols available. The DMP-41 is easy to live with, adhering to FCC Class B requirements. UL listing pending.

POINT . . . Minutely defined step size and high-resolution logic—combined with robust drives and optimized pen ballistics enable you to create plots of high precision and surpassing quality.

POINT . . . The Houston Instrument DMP-41 is one of your most cost effective considerations.*

For the name, address and phone number of your nearest representative, write Houston Instrument, P.O. Box 15720, Austin, TX 78761. Phone 512-835-0900, or 800-531-5205 if outside Texas. In Europe contact Houston Instrument Belgium NV., Rochesterlaan 6, 8240 Gistel, Belgium. Tel 059-27-74-45, tlx 846-81399.

houston instrument

Houston instrument 25th Anniversary 1959-1984

*suggested US retail $2,995

CIRCLE 3

4P009B

UP FRONT

System will attack fault-tolerant price barrier

Founded less than a year ago, EnMasse Computer Corp (Acton, Mass) is in the late stages of developing a 32-bit Unix-based multiprocessor system with fault-tolerant features. The system is currently scheduled for introduction during the first quarter of 1985 and is list priced below $50,000 for a basic configuration that will support up to 32 users. Customers will be able to expand on the basic hardware to build a single system that can support up to 768 user terminals while handling multiple tasks simultaneously, all with no sacrifice in processing speed. Incremental cost for adding each new user is claimed to be reduced to as low as $500. The computer runs under Unix System V, the emerging *de facto* standard operating system for program portability. EnMasse Link architecture creates a high speed network of file processors and distributed application processors. The system will have gateways to other EnMasse computers, IBM and DEC hosts, and IBM PC and compatible workstations. The system is designed for distributed online transactions, office, communications, and other multi-user applications.—*J.H.*

Low cost PC video imaging system performs at less than 1 s

A single-board, high performance frame grabber-based subsystem to be released within the next month or two by Data Translation (Marlboro, Mass) fits into any expansion slot of an IBM PC or PC/XT and provides realtime video digitization and display. The DT2803 hardware will sell for about $1500 and operates with realtime video I/O software called Videolab. Comparison benchmark tests for the board alone in a PC indicate a required time of 10 min for a 3 x 3 convolution pixel operation. With a math coprocessor board in place, the operation required only 50 s. However, the frame grabber board is designed to operate optionally with an array processor board, to be supplied by Sky Computers (Lowell, Mass). With that board, the operation requires 1 s or less.—*S.F.S.*

Automated design processes—heading toward complete integration?

In the CAD/CAE/CAM environment, some equipment offers only schematic capture, while others will not only help you with the design, but will carry the process all the way through to production. The most apparent trend, as noted at the recent Design Automation Conference, is that the industry is heading for total integration of the design process, from chip to system level. Even such ATE firms as GenRad (Concord, Mass) and Teradyne (Boston, Mass) are becoming links in the design chain and offer entries that provide logic design simulation throughout an entire design evolution, as well as the test generation to improve all testing efficiencies. These procedures include program generation testing, prototype testing, and digital IC manufacturing testing. While integration seemed to be the main theme of the conference, the wide range of available hardware and software indicates that everyone can become involved in design automation. The low end process, usually running on the IBM PC, offers the potential to speed the design cycle at minimum cost. At the high end, Control Data Corp provides complex software that includes an integrated data base and engineering data library. The power behind this approach is that it allows simulation of complete systems rather than of individual components.—*M.B.*

UP FRONT

Design tools simulate entire system

Design tools that Control Data (Minneapolis, Minn) engineers used internally to design ICs for the CYBER 200 series supercomputers are now available as commercial products. The company's Modular Integrated Design Automation System (MIDAS) consisting of logic simulation, fault simulation, and layout has been developed and continually enhanced by Control Data for the past 10 years. MIDAS operates on the company's CYBER computers and uses the Cybernet data services network to transfer integrated data bases and engineering data beween stations. Designers can thus simulate complete systems rather than being limited to simulating individual components. Control Data engineers have designed boards with 120,000 logic gates using MIDAS, a hundred-fold increase over current densities. MIDAS' shared data base allows the designer to make numerous design and layout changes before a board is committed to fabrication. Control Data claims that its engineers routinely perform system-level simulation following completion of the physical design steps for system complexities up to 300,000 gates. A minimum, fully configured system based on the CYBER 810, with all the necessary hardware and MIDAS software applications, costs $500,000.—*N.M.*

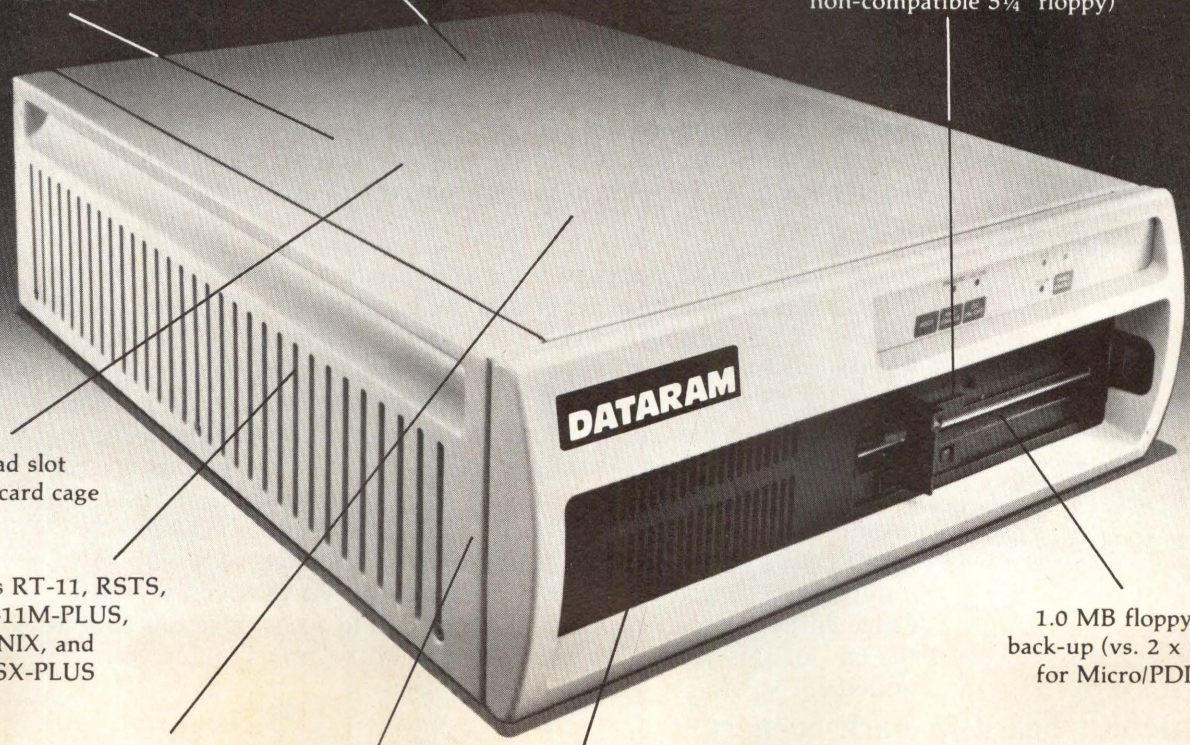
NCC '84 melts away in Vegas

Although the American Federation of Information Processing Societies was eager to predict a large attendance in its promotion for NCC '84, no one at that organization will now release even an estimate of the final attendance figures. Preconference estimates from AFIPS, the central organization among the several sponsors of all National Computer Conferences, were for a minimum of 100,000 attendees. Whether it was the expected oppressive heat of Las Vegas in July or any of a seemingly endless number of reasons offered by those persons who did attend, that original figure was not even nearly met. Estimates from among a number of press and exhibitor personnel who have been to many National, Spring Joint, and Fall Joint Computer Conferences agreed that the final attendance was probably not over 60,000—and even that may be a generous figure. Additionally, last-minute cancellations left several booth spaces unfilled, a true novelty for NCC in the past several years. Whatever the true reasons were for this year's ultralow attendance, NCC '85 will be in Chicago—and although it again will occur in July, no one presumably will be able to logically blame the heat if people or exhibitors stay away—*S.F.S.*

Semicustom IC design—a movable service

In what amounts to taking the semicustom IC design business to the customer instead of having the customer seek out a semicustom chip vendor, RCA Solid State (Somerville, NJ) has signed an agreement with distributor Hamilton-Avnet (Culver City, Calif) to both design and sell RCA's COSMOS gate array and standard cell ICs. H-A plans to open 12 design centers across the United States over the next year to assist users in the logic design of semicustom chips. Customers will be able to implement the completed logic design and debug the chip at workstations set up at the design centers and then forward the MIMIC data base to RCA's manufacturing facilities for fabrication. In this way, the customer will be spared the necessity of learning the rest of the design automation tools and will only need to learn the MIMIC software. This market is expected to reach a $2 billion volume for gate arrays and standard cell ICs by 1987, according to Integrated Circuit Engineering Corp, a research firm based in Phoenix, Ariz.—*N.M.*

Like DEC's.



$8,845 system price*

256 KB minimum... up to 4 MB!

Media and software compatibility with DEC's RX02 8" floppy (vs. Micro/PDP-11's non-compatible 5¼" floppy)

8-quad slot Q-BUS card cage

Supports RT-11, RSTS, RSX-11M-PLUS, UNIX, and TSX-PLUS

Two fans in card cage area (vs. one in Micro/PDP-11)

RL02-compatible 5¼" Winchester disk; 10 MB, 20 MB, or 40 MB capability

Cartridge tape capability

1.0 MB floppy disk back-up (vs. 2 x 400 KB for Micro/PDP-11)

Only better.

You can buy DEC's Micro/PDP-11 with its impressive array of features...or you can get Dataram's A22 —an LSI-11/23 based minicomputer that gives you a whole lot more...for a lot less dollars! Like an 8" RX02-compatible floppy. 40 MB 5¼" Winchester and ¼" cartridge tape capability. And two fans that provide push-pull air flow in the card cage area.

For more information, forward this coupon to us, or, for faster response, call (609) 799-0071.

☐ Send information. ☐ Contact me immediately.

Name

Company

Address

City State Zip Phone

Dataram Corporation, Princeton Road, Cranbury, NJ 08512

*$8,845 is single-quantity domestic price for A22 with LSI-11/23, 256 KB, 10 MB Winchester and RX02-compatible 8" floppy.

DEC, LSI-11, Micro/PDP, PDP, RSTS, RSX, and RT-11 are trademarks of Digital Equipment Corporation.
TSX-PLUS is a trademark of s&h computer systems, inc.
UNIX is a trademark of Bell Laboratories.

CIRCLE 4

DATARAM

Dataram Corporation ☐ Princeton Road ☐ Cranbury, New Jersey 08512 ☐ Tel: 609-799-0071 ☐ TWX: 510-685-2542

UP FRONT

AI development tools available for the professional personal computer

An expert system development package using the Professional Computer from Texas Instruments (Dallas, Tex) incorporates all necessary features for a designer to prototype and develop sophisticated commercial applications. The Personal Consultant development tools are aimed at users who are interested in developing rule-based expert systems at their desktops without the need to learn the popular AI Lisp language. TI researchers have coded their program using a similar data file format to the EMYCIN developed at Stanford University. As opposed to programming MYCIN in Lisp, users of the Personal Consultant can write their rules in a Basic-like rule specification language. The tools will be available in the last quarter of 1984 for $3000 a copy. A 3-day training course is also available for $1500. In addition, the package can be obtained bundled with the Professional Computer for $15,500. Creating end-user knowledge bases from the tools can be done for a one-time license fee of $25,000 and a royalty fee of $25 for each application copy.—*N.M.*

PDI handles data transfer PDQ

A new twist in the standard interface situation, the Parallel Device Interface (PDI) promises to bridge the gap between high and low level devices. To accomplish this, PDI, a level 2 definition of the Intelligent Peripheral Interface (IPI) specification, will support data transfers occurring at rates between 15 and 80 Mbits/s. Promoted by Control Data Corp (Minneapolis, Minn) and supported by Emulex (Costa Mesa, Calif), the specification provides for maximum transfer rate over a maximum cable length, simplifies controller design, and promotes future compatibility. In addition to supporting the specification, Emulex is managing the development of a chip to implement PDI and thus ensure compatibility to controllers. PDI is expected to see red line print this month and to achieve blue line by the end of this year. First products should be forthcoming by mid-1985.—*P.K.*

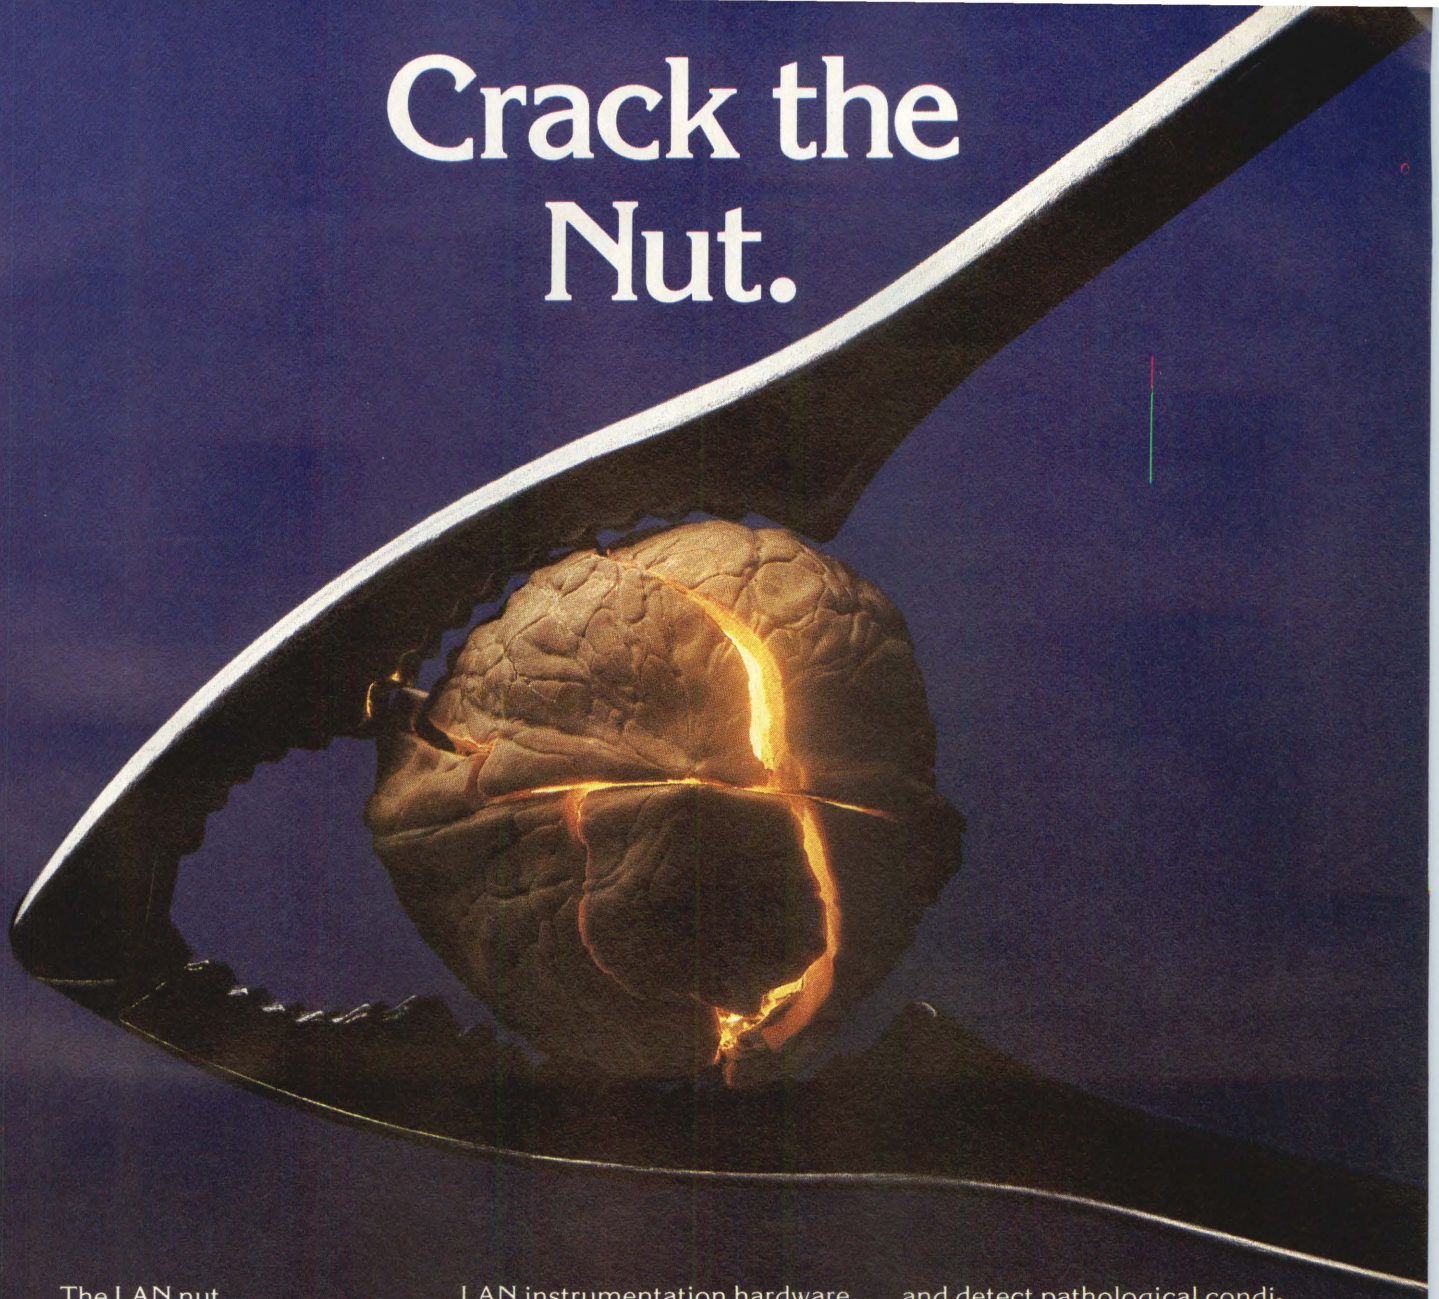
Scuttlebutt heard in dark corners

Rumors and general data wafting around the environs of the recent NCC '84 hint that a major manufacturer is readying a featherweight portable computer—so light, in fact, that it needs no handle. Details are sparse, save that it is said to incorporate an LCD, two sub-5¼-in. disk drives, and the obligatory keyboard. There's also no handle as to price and delivery. Keep tuned—*J.H.*

Unix starts to appear in smaller packages

Efforts to provide real Unix environments to an ever-increasing community of users are meeting with some success, thanks to some of the newer microprocessors and advances in memory technology. At the same time, this push is spurring developments in other areas. Unix means disk storage, so developments in half-height, large capacity Winchester drives are rising to meet the challenge. Advanced Storage Technology and Qume Corp (both of San Jose, Calif) have announced half-height 5¼-in. Winchesters using plated media. The Qume line ranges in capacity from 13 to 40 Mbytes, and the AST from 61 to 220 Mbytes. Processors such as the Intel iAPX 256 and the Motorola 68010, which support virtual memory, are making it possible to build a range of multi-user, Unix-based systems that can support different numbers of users but that aim for a per-user cost of about $2500. With the small, high capacity drives becoming available, these systems should appear in even smaller packages as well.—*T.W.*

Crack the Nut.



The LAN nut.

When you're developing a LAN-based computer system, you're faced with a tough nut to crack. You need to know what's going on in your Local Area Network system. Many computers are carrying on many conversations. Simultaneously. And at speeds 10,000 times faster than traditional data communications. How do you test it? How do you debug it? How does the system really perform?

Excelan can help you crack that nut with the Nutcracker™, the world's only comprehensive analyzer/simulator for Ethernet systems. Packaged as an integrated workstation, the Nutcracker provides advanced LAN instrumentation hardware capable of making real-time diagnostic decisions at 10 million bits per second. Also included is an 8086-based CPU, about 1 MB of RAM, a 20 MB disk, a 600 KB floppy, keyboard, 12″ CRT and 100 cps external printer, and complete menu-driven software that brings the power of the Nutcracker to you.

With the Nutcracker connected to the Ethernet cable, you're in control. You can create and detect pathological conditions for shake-down testing. Powerful filtering facilities allow you to extract precise packet substreams for triggering and tracing functions to find those elusive bugs. You can generate traffic for simulation of various load conditions. And high resolution time-stamping means you can track system performance against this regulated traffic load.

Increase your productivity, speed development time, and get your network system up and running — with confidence in it's performance and integrity.

Join computer industry leaders. Crack the LAN nut with the Excelan Nutcracker.

EXCELAN

"Excellence in Local Network Technology"

2180 Fortune Drive
San Jose, California 95131
(408) 945-9526 TELEX 176610

CIRCLE 5

Nutcracker is a trademark of Excelan, Inc. Copyright Excelan, Inc., 1984. All rights reserved.

HERE'S A TANDON DISK DRIVE WITHOUT TANDON-BUILT PARTS.

As you can see, the Tandon theory of vertical integration doesn't leave much room for outside manufacturers. That's precisely what we liked about it from the start. And we've been right about our hunch.

We're the leading manufacturer of microcomputer random access disk drives in the world. We didn't get here by following in the paths of our competitors.

Instead of piecing together someone else's parts, we manufacture our own. We know every detail about every part in all our drives. From manufacturing the button head for the recording head assembly, to producing our own plated media, we're involved with it all. As a consequence, Tandon-built parts comprise 80% of the cost of our drives. There's a lot to be said for that.

Quality, for one thing, is a lot easier to control. The toughest Tandon specifications are met every step of the way. Revisions and improvements can be implemented when they're needed, where they're most effective, resulting in higher reliability.

That tight control makes it easier for us to keep our costs down, and manufacturing yields high. Per unit savings mean the lowest overall prices available. Which is obviously to your advantage.

And since we're a supplier as well as a manufacturer, we'll get your order to you quickly. We'll always deliver the drives you need when you need them.

We've recently made a move to expand our vertical integration even further. Our new $5 million facility produces plated disks for use in all of our Winchester drives. In keeping with Tandon quality, cost-effectiveness, and advanced technology.

All this is what we're calling vertical integration. A strategy we initiated before many of our competitors even existed. But what it comes down to for you is a reliable, high-performance disk drive at the lowest of possible costs.

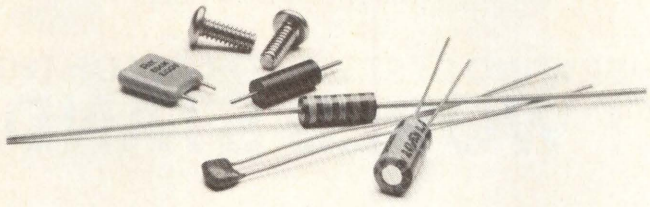
And, all theory aside, you just won't find a better combination than that.

Tandon

THE DRIVING FORCE BEHIND THE SMALL COMPUTER INDUSTRY.

Tandon Corporation, 20320 Prairie, Chatsworth, CA 91311. (818) 993-6644, TWX: 910-494-1721, Telex: 194794. Regional Sales Offices: Boston (617) 938-1916 • New York (201) 851-2322 • Atlanta (404) 934-0620 • Chicago (312) 530-7401 • Dallas (214) 423-6260 • Irvine (714) 669-9622 • Santa Clara (408) 727-4545 • Frankfurt, West Germany 6107-2091, Telex: 411547 • London, England (0734) 664-676, Telex: 848411. Distributors: Hall-Mark, Kierulff, Schweber.

CIRCLE 6



UNIX® IS A DINOSAUR
CP/M® & MS-DOS™ ARE TOYS

MULTI SOLUTIONS PRESENTS

THE WORLD'S FIRST 4th GENERATION OPERATING SYSTEM

A SERIOUS OPERATING SYSTEM **S1**™ FOR TODAY AND TOMORROW

- PORTABLE
- MODULAR
- MULTIUSER
- MULTITASKING
- MULTI PROCESSING
- PARALLEL PROCESSING
- 64 CHARACTER NAMES
- 3 COMMAND PROCESSORS
- REAL TIME
- NETWORKING
- DISTRIBUTED PROCESSING
- HIERARCHICAL DIRECTORIES
- KEYED FILES
- ISAM
- VSAM
- B-tree
- RECORD LOCKING
- UNIX SOURCE COMPATIBLE
- WINDOWING
- BIT MAPPED DISPLAYS
- FULL SCREEN MANAGEMENT
- FULL SCREEN EDITING
- FULL MEMORY MANAGEMENT
- VIRTUAL MEMORY
- SEMAPHORES & LOCKS
- EXTENSIVE UTILITIES
- AND MUCH, MUCH MORE

S1 IS THE ONLY OPERATING SYSTEM WORTHY OF THE TITLE: "THE NEXT WORLD STANDARD."

ONLY S1 DOES IT ALL.
NO OTHER OPERATING SYSTEM COMES CLOSE. CUTS DEVELOPMENT TIME FROM MAN YEARS TO MAN MONTHS.

MULTI SOLUTIONS INC.™
Multi Solutions, Inc.
123 Franklin Corner Rd.
Lawrenceville, N.J. 08648
609-896-4100 Telex: 821073

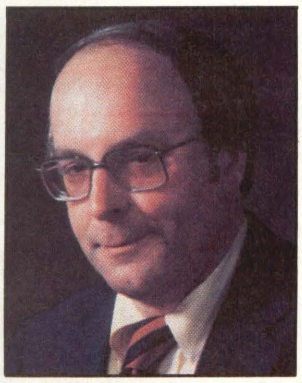
Please send for our **FREE S1 Book or Call 609-896-4100**

IN TIME, ONLY THE BEST WILL SURVIVE: S1

*Reg. Trademarks: CP/M of Digital Research; Unix of Bell Laboratories Trademarks: MS-DOS of Microsoft, Inc.

CIRCLE 7

SHOWTIME IN LAS VEGAS



As I write this, it is July and I am at the National Computer Conference in Las Vegas. The people at our office in Massachusetts tell me my editorial is late and that I have to rush it to them today. But, the temperature hit 110 degrees today, and my usual problem of technostress has been replaced by a new problem—heat stress.

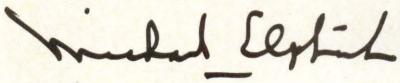
This town appears to be controlled by its cab fleet owners. Because parking facilities are inadequate for major conventions, conventioneers are forced to the streets in search of cabs. This week, people waited an hour or more for cabs, and then were jammed, four at a time, into cabs without air conditioning. As if this weren't enough, some of the drivers had the nerve to collect the full fare from each rider. Any city that is serious about attracting convention business should have corrected its transportation problems long ago. But then, as we heard today from some of the casino owners, Las Vegas doesn't like conventioneers—especially those from our industry—because they don't gamble enough. Did the casino people honestly believe that otherwise-intelligent engineers came to Las Vegas in July to have fun? How could they really believe that entrepreneurs and salespeople in the fiercely competitive computer industry would need the extra risk and excitement of the crap tables to spice up their lives?

Besides, physical survival in Las Vegas can be quite a gamble. In fact, if you have ever wondered what happens if the power fails here, I can now give you a first-hand account. When the subject is under discussion, the standard response is "Don't worry, they have backup power." Well, after 60 mph winds downed some city power lines, I was able to see "Plan B" in action.

Yes, the hotels do have backup power; but that doesn't begin to supply the total need. So, with a characteristically inverted set of priorities, the gaming tables and slot machines took precedence over "less important" things like air conditioning and elevators.

I think it's obvious that we don't need this city and this city doesn't need us. So, why are we here? Unfortunately, Las Vegas is one of the few cities with convention and hotel facilities extensive enough to accommodate a show the size of NCC. An added problem is that this year NCC is competing with other events such as the Olympic games and the national political conventions. As a result, Las Vegas grabbed most of the big computer shows this year. Thus, we were here in March for Interface, we are here now for NCC, and some of us will be here again in November for Comdex.

There are no easy solutions to the problems posed by trade shows and conferences. There are too many of them and they are becoming too large. Ironically, some of the better organized ones are already being strangled by their own success. We can't expect the organizers or exhibitors to voluntarily cut back on their activities. That would be conceding victory to their competitors. The shakeout will be precipitated only by a market boycott—and we are the market. We have to vote with our feet and start attending only those conferences that are vitally important to us—and at the best times and in those locations where the conditions are most tolerable. Therefore, I hereby serve notice that, although I will be attending Wescon in Anaheim, I will not be attending Comdex in Las Vegas. Other editors will represent our magazine, but I won't be one of them—we are cutting back on our coverage.



Michael Elphick
Editor in Chief

YOU WON'T FIND MANY SHUGART DRIVES AROUND THESE PARTS.



You find Shugart new generation drives where they were built to be. In all kinds of systems. Under all kinds of conditions. Running. And running and running.

Setting a world standard for quality. A standard that doesn't allow room for failure. Because you don't.

To insure the reliability of your system, we begin with our suppliers. They go through a lot. Learning and qualifying, even building drives themselves.

We also involve our manufacturing engineers. And our quality control experts. From the very beginning. The design stage.

With their full collaboration, every new drive is designed for top quality. Unbeatable reliability. And ease of assembly. Using the fewest number of parts possible.

In other words, the same uncompromising quality you design into your system is designed and built into every new Shugart drive. Right from the start. So you can count on consistent reliability. For example, 20,000 hours MTBF for new generation Winchesters.

Of course, every Shugart drive is subject to the most stringent testing and inspection at every stage. But inspection doesn't insure quality. At Shugart, we believe there's only one way to insure a reliable product. By eliminating the cause of failure from the product design.

At Shugart, quality is designed in. Not inspected in.

Let us prove it. Write for our white paper on quality to Shugart Corporation, 475 Oakmead Parkway, M/S 3-5, Sunnyvale, CA 94086. Or call, (408) 737-4360.

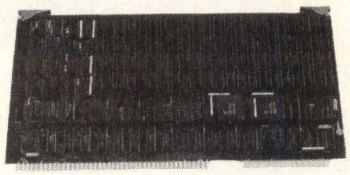
Shugart

Right from the start.

Milpitas, CA (408) 263-2600; Costa Mesa, CA (714) 979-1935; Thousand Oaks, CA (805) 496-5388; Rochester, NY (716) 235-7190, Minneapolis, MN (612) 546-4411; Richardson, TX (214) 234-3568; Framingham, MA (617) 879-1700; Saddle Brook, NJ (201) 368-8445; Smyrna, GA (404) 436-0953; Markham, ONT (416) 475-2655; Paris, France (3) 946-42-66; Munich, West Germany (089) 786-021; London, U.K. (44) 4862-27272; Wanchai, Hong Kong (852) 5-733307. © 1984 Shugart Corporation

CIRCLE 8

UNFORGETABLE MEMORIES
MULTIBUS — QBUS

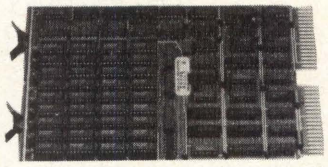


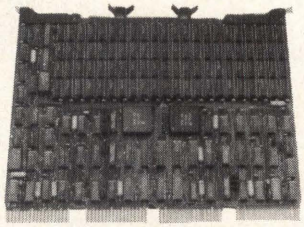
MULTIBUS MEMORY
512KB TO 2MB EDC

- Error Detecting and Correcting (EDC)
- Faster Access Times

SINGLE QTY. PRICE

	Without/EDC	W/EDC
512KB	895.00	1495.00
2MB	4775.00	5775.00





Q-BUS MEMORY
256KB TO 4MB

- Battery Backup Mode
- Block Mode DMA, ECC or Parity
- Works with LSI-11, J-11, MicroVAX

SINGLE QTY. PRICE

1MB 1975.00
256KB 525.00

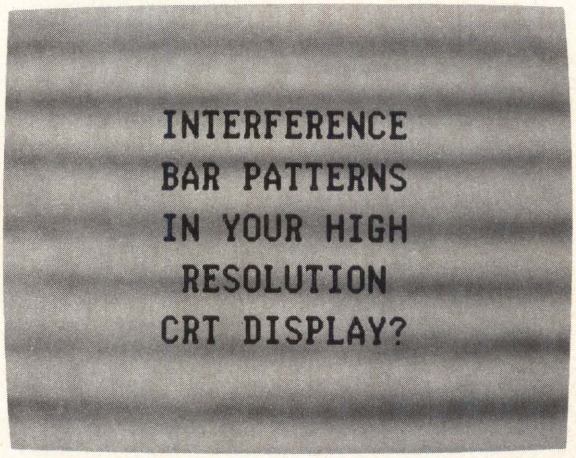
"OFFERING QUALITY WITH AFFORDABLE PRICING"

Chrislin Industries, Inc.

31352 Via Colinas • Westlake Village, CA 91362 • 818-991-2254
TWX 910-494-1253 (CHRISLIN WKVG)

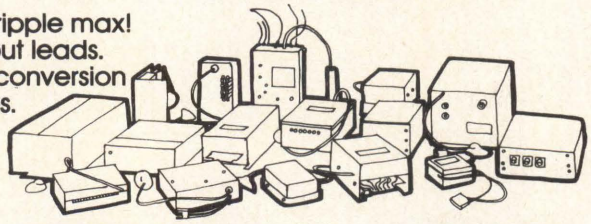
J-11, LSI-11, MicroVAX, Q-BUS are trademarks of Digital Equipment Corporation. MULTIBUS is a trademark of Intel Corporation.

CIRCLE 9



KELTRON has the answers in High Voltage Power Supplies for the color or monochrome OEM...

- 1V anode ripple max!
- Filtered input leads.
- Choice of conversion frequencies.



Custom design HVPS for the OEM
KELTRON CORPORATION
225 Crescent Street, Waltham, MA 02154 • (617) 894-8700

See what Keltron has to offer!

CIRCLE 44

COMPUTER DESIGN

The PennWell Building,
Littleton, MA 01460, Tel: (617) 486-9501
Editorial/Executive Offices

Editor in Chief, Michael S. Elphick

Managing Editor, Sydney F. Shapiro
Senior Editor, John Bond
Senior Editor, Peg Killmon
Special Features Editor, James W. Hughes
Senior Associate Editors,
Malinda E. Banash, Deb Highberger
Associate Editor, Jack Vaughan
Assistant Managing Editor,
Leslie Ann Wheeler
Copy Editors,
Helen McElwee, Leah A. Rappaport, Nancy E. Purcell, Jane E. Shattuck
Editorial Assistants, Julia E. Cote, Lisa M. Stephens
New York Field Office:
230 Park Ave, Suite 907
New York, NY 10169, Tel: (212) 986-4310
Senior Editor, Nicolas Mokhoff
Special Features Editor, Harvey J. Hindin
Western Field Office:
540 Weddell Dr, Suite 8
Sunnyvale, CA 94089, Tel: (408) 745-0715
West Coast Managing Editor,
Tom Williams
Senior Editor, Bill W. Furlow
Field Editor, Joseph A. Aseo
Editorial Assistant, Robin Mock

Production Director, Linda M. Wright
Production Manager, Philip Korn
Art Director, Lou Ann Morin
Technical Art, Designline
Ad Traffic Coordinator, Debra Friberg
Printing Services, Padraic Wagoner

Marketing Director, Robert A. Billhimer
Circulation Director,
Robert P. Dromgoole
Promotion Director, Steve Fedor
Marketing Services/PR Manager,
Linda G. Clark

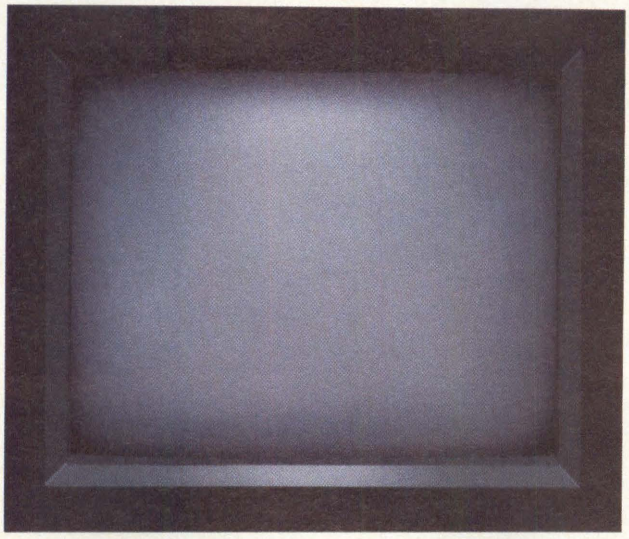
Publisher, Frederic H. Landmann

PennWell

Advanced Technology Group
119 Russell St, Littleton, MA 01460
Tel: (617) 486-9501

H. Mason Fackert, *Group Vice President*
Saul B. Dinman, *Editorial Director*
John M. Abernathy, *MIS/DP Director*
Patricia M. Armstrong,
Administrative Services Manager

Before TEMPLATE®, graphics software that ran on mainframes and minis and micros all looked like this.



Which is to say, nonexistent. Some manufacturers make graphics software packages for mainframes. Others for minis. Still others make it for micros. But no one made high-level graphics software that ran on all three.

TEMPLATE just changed all that. By becoming the *only* high-level graphics software available on micros, minis and mainframes. And it took our extensive experience in graphics software to do it.

Now all computer-using design engineers and scientists can utilize the industry's finest software. And bring mainframe applications right to the bench. Or vice versa. Which means TEMPLATE's device-intelligence and computer-independence is even further enhanced. And you get the graphics functionality for mechanical and electrical CAD, scientific analysis, seismic work, VLSI, and molecular modeling that puts TEMPLATE in a class by itself. In any environment, whether it's batch or interactive, 2D or 3D.

TEMPLATE features table-driven architecture, 3D software display lists, metafile capability over 250 user-callable FORTRAN routines, workstation model, post processing capability, run-time selection, and complete support functions.

So when you're looking for graphics software that'll run on *all* your computers, call Megatek.

And find out all about TEMPLATE.

The product that just gave micros, minis, and mainframes a new computer graphics image.

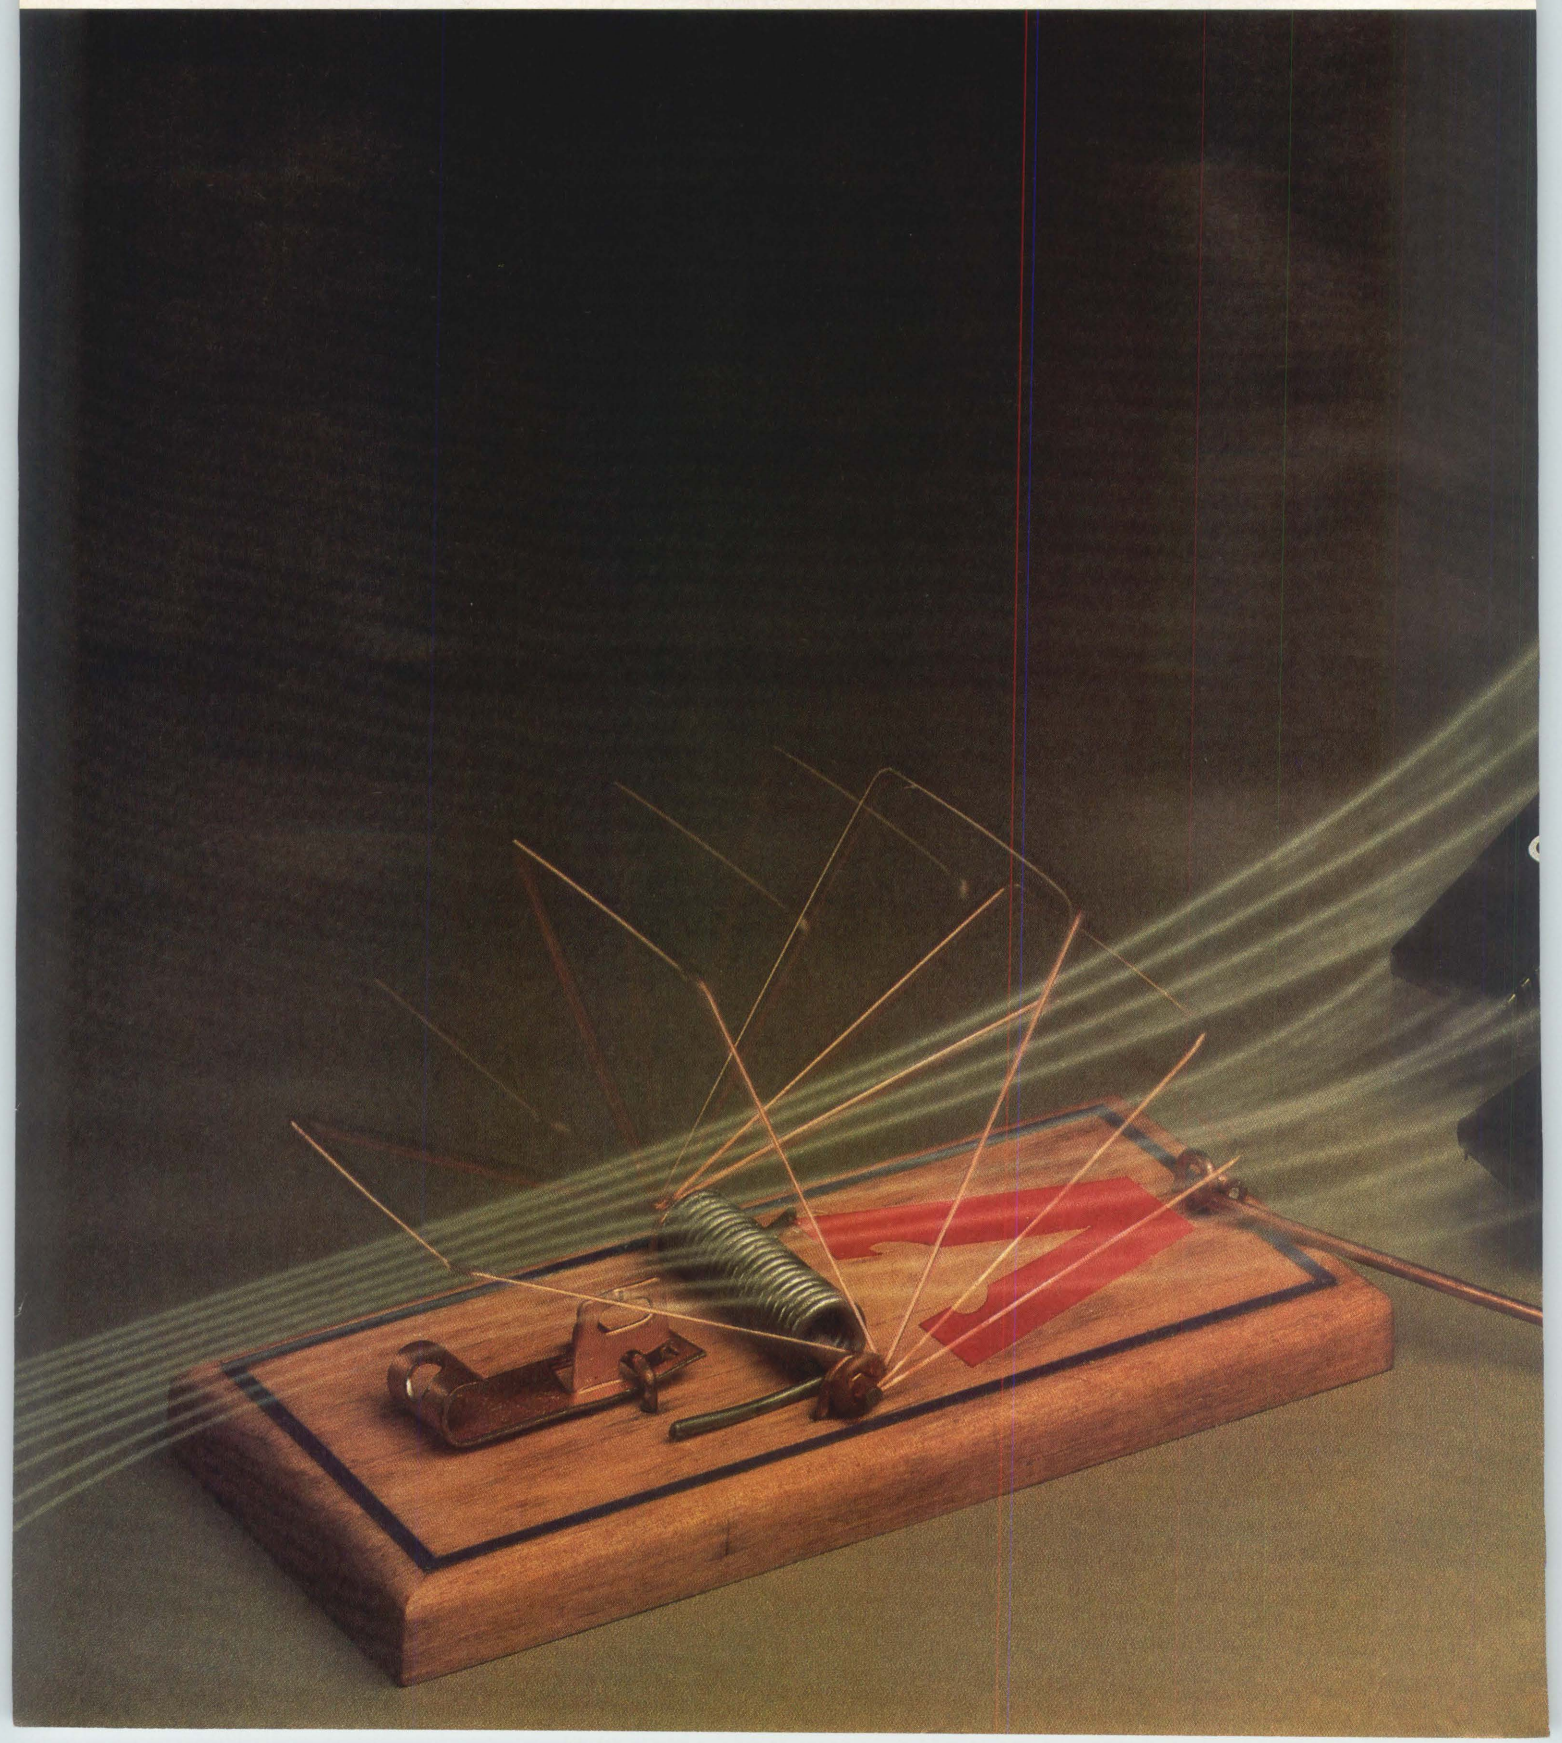
MEGATEK CORPORATION
® A UNITED TELECOM COMPANY
Making History out of State-of-the-Art

World Headquarters • 9645 Scranton Road • San Diego, CA. 92121 • 619/455-5590 • TWX: 910-337-1270

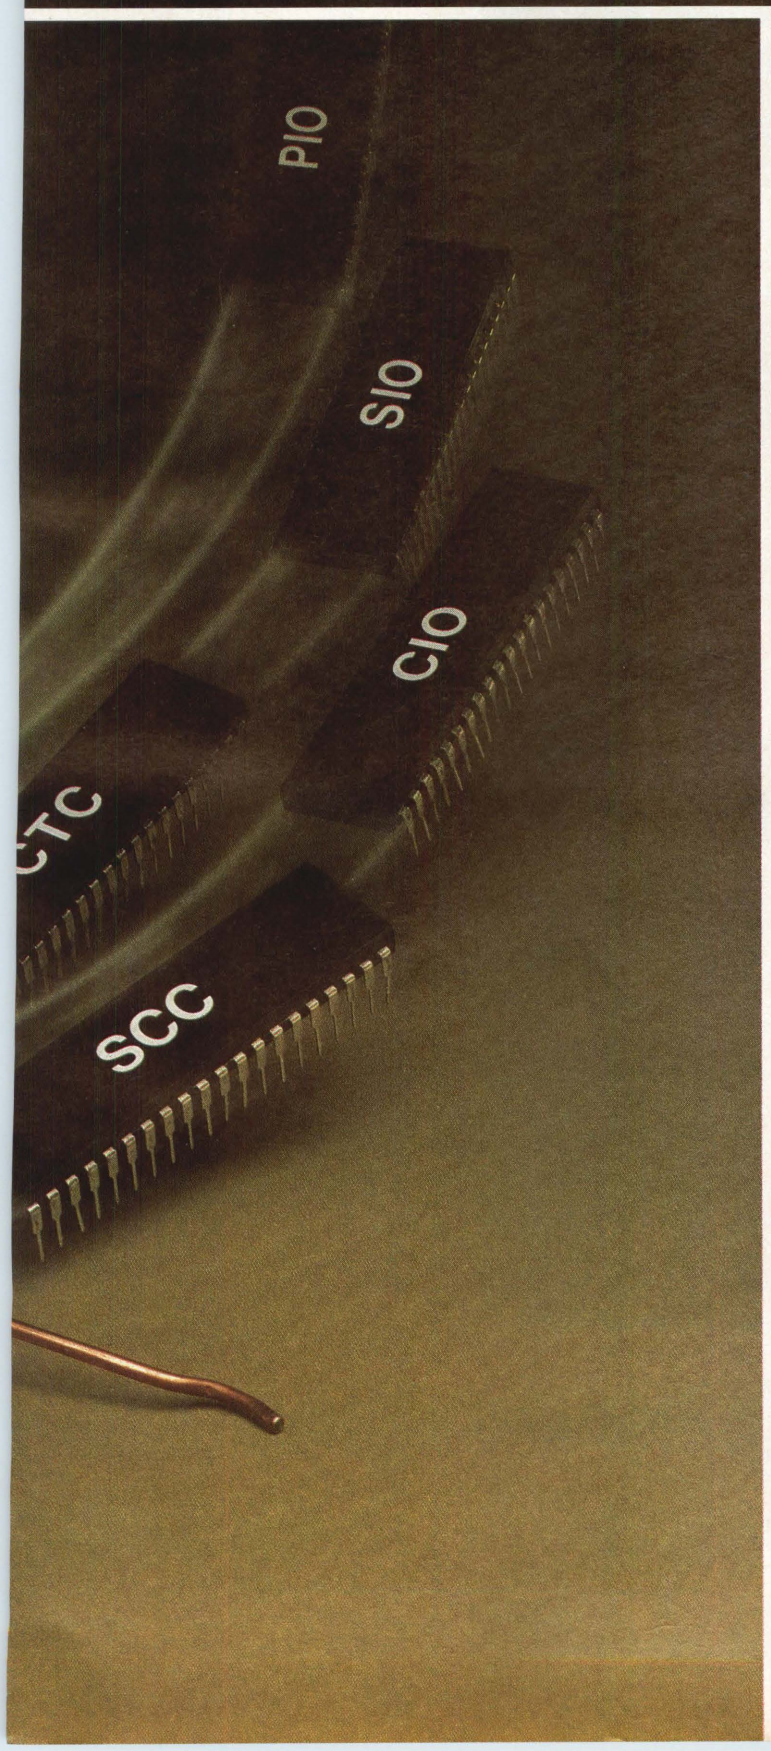
TEMPLATE is a registered trademark of MEGATEK CORPORATION.

CIRCLE 11

You'll never get trapped into dead-end designs with



Zilog's high-speed Z80B CPU's and Peripherals.



Here's more proof that nobody does more to extend the life of your 8-bit designs than Zilog. Because now you can increase 8-bit Z80® performance up to 6 MHz with the high-speed Z80B CPU and its family of peripherals. You can join the hundreds of design engineers that have already tested this claim with winning results. Or wonder...

Is there something here they know that you don't?

Like the fact that the Z80B CPU has the same 158 instruction set and the elegant registers and interrupts that you're used to working with, but runs them 50 percent faster than the Z80A chip?

That the Z80B processor is completely software compatible with the rest of the Z80 family, permitting you to upgrade to higher performance without getting trapped into software redevelopment? That software compatibility also means you can use the Z80B device in co-processing and/or multi-processing environments along side our other Z80 processors?

And consider the fact that you can surround our Z80B CPU's with a complete family of Z8400 and Z8500 peripherals and really boost system performance. They help you keep your parts and space requirements to a minimum and increase system throughput because we build more functions into every device. The peripherals include a PIO, a CTC, an SIO, an SCC, an FIO, an FIFO, a CIO, and a UPC.

For complete specifications and applications data on the Z80B and peripherals, fill out the coupon and mail to: Zilog, Inc., Components Tech. Publications, 1315 Dell Avenue, MS C2-6, Campbell, CA 95008. Or call our TOLL FREE Literature Hot Line at 800-272-6560. For information on Zilog's other components, call (408) 370-8000.

Z80 is a registered trademark of Zilog, Inc.

☐ I'd like more information at this time.
☐ Please have a salesman contact me.

Name_____

Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Phone_____/_____

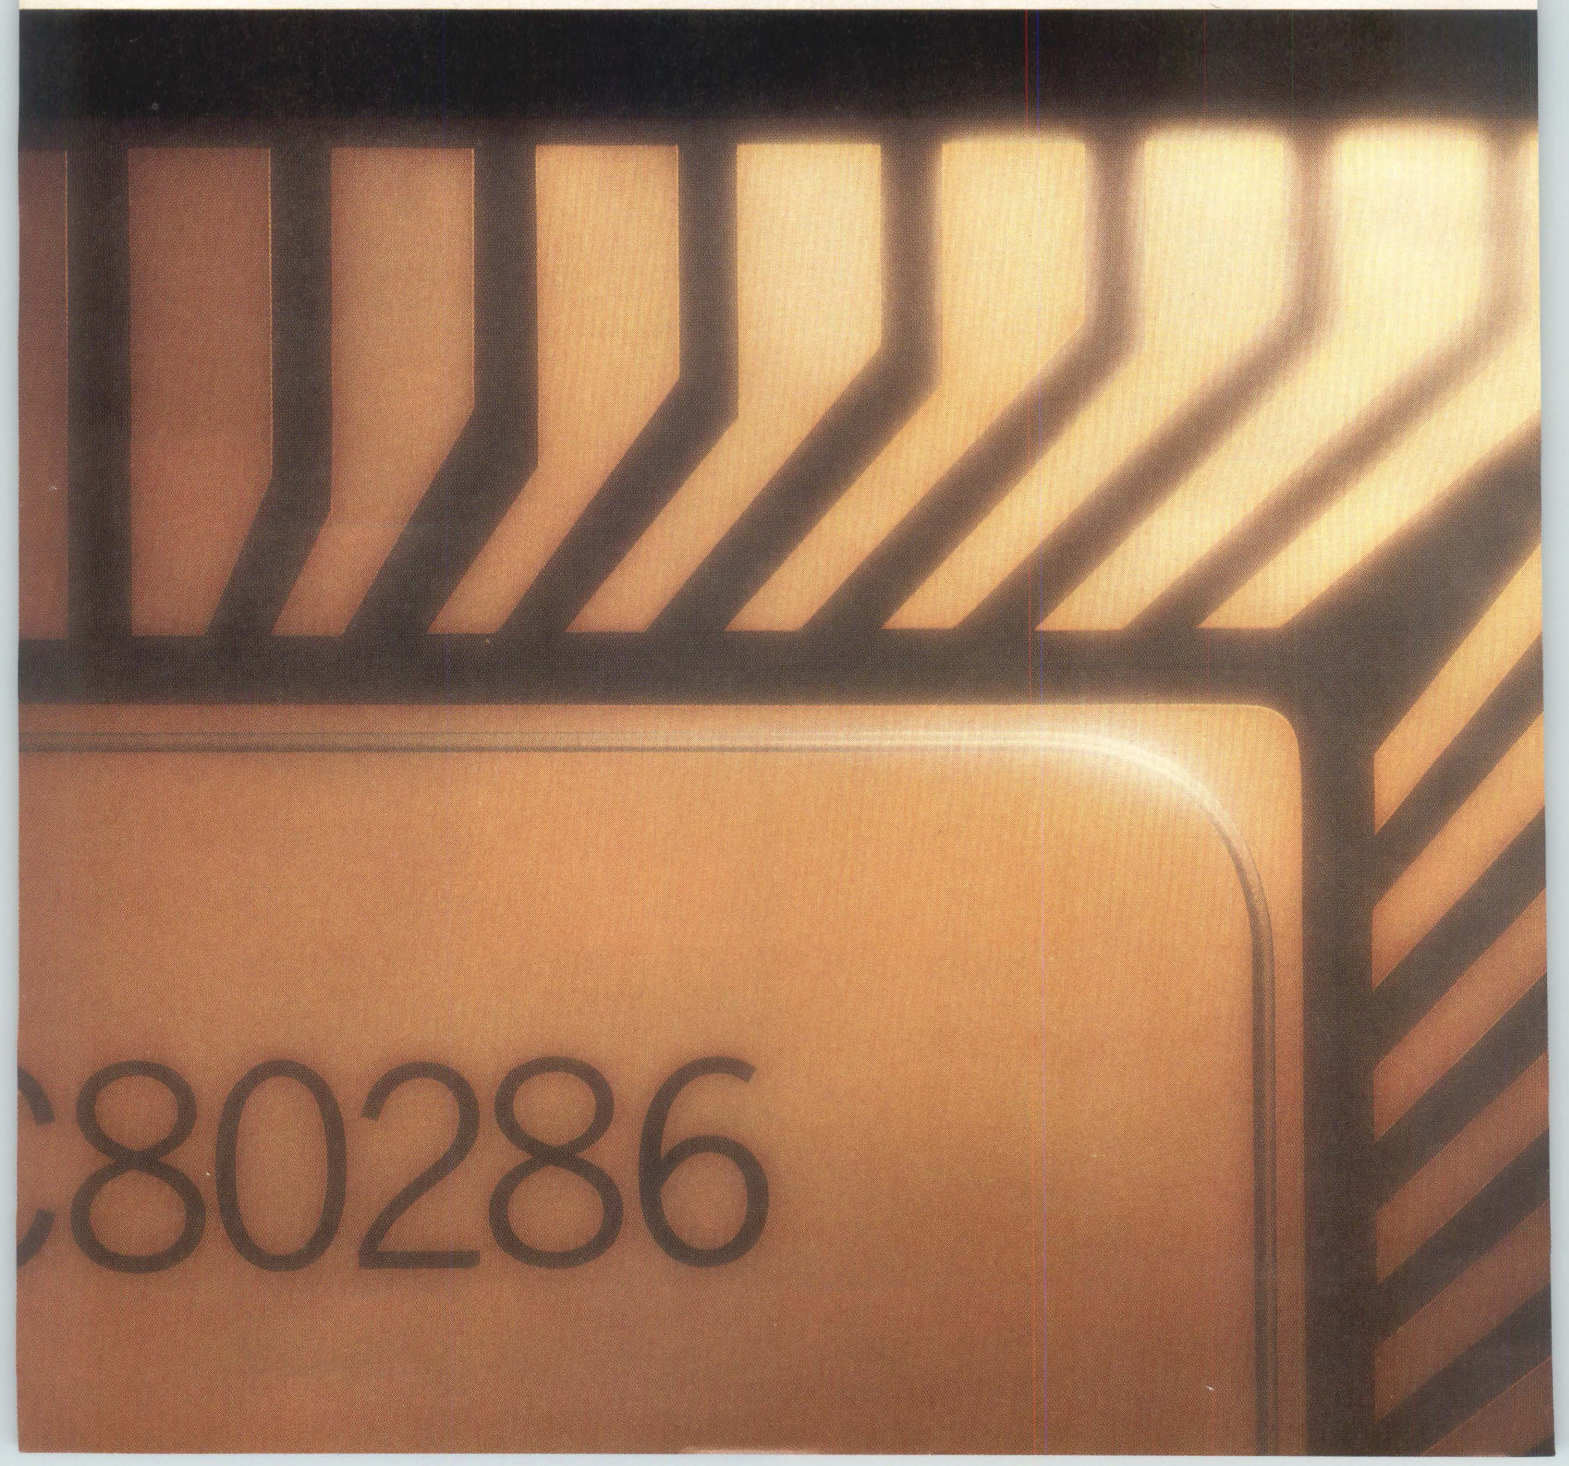
Z80B/Peripherals CD-8/84

Zilog
an affiliate of
EXXON Corporation

Pioneering the Microworld

CIRCLE 12

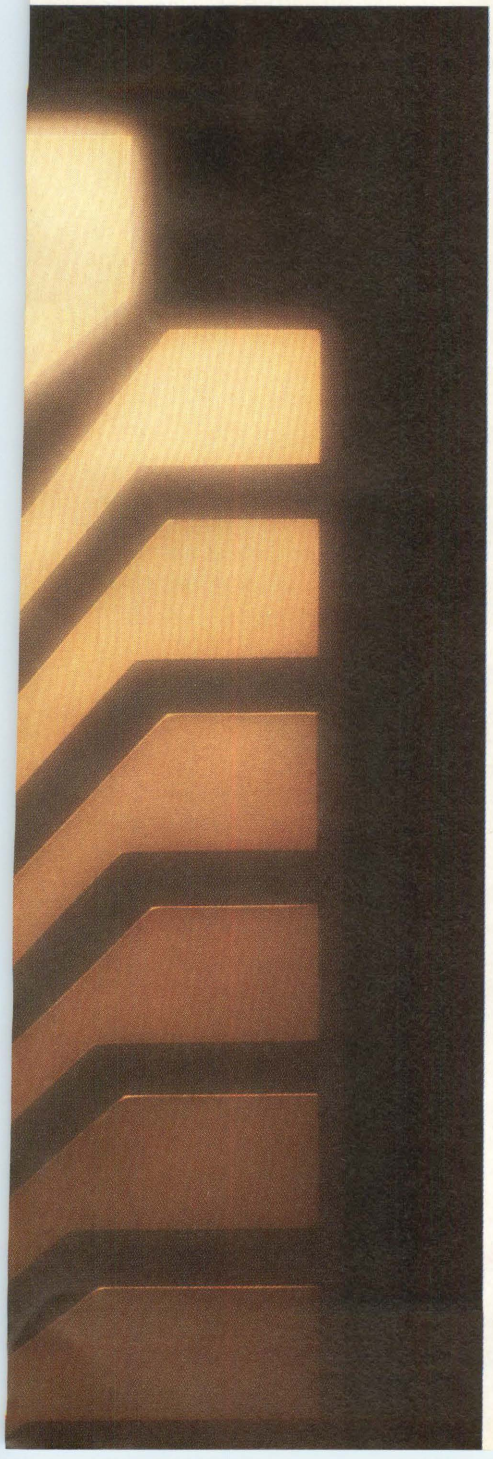
MICROSOFT STRIKES GOLD.



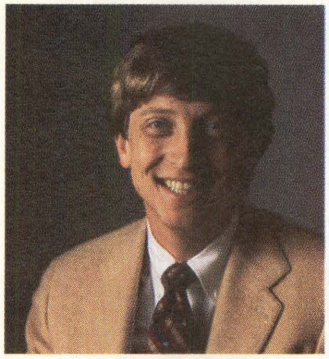
This is what paydirt looks like when it's turned into silicon.

Intel's iAPX 80286 microprocessor.

Together, Microsoft and Intel used this gem to create another: the XENIX* 286 Operating System. The most powerful implementation of UNIX* on the market today for any microprocessor. And the most reliable, functional, commercial enhancement of the UNIX operating system ever created.



The XENIX 286 Operating System includes UNIX enhancements available from both AT&T™ and U.C. Berkeley. But Microsoft has also added loads of other commercial enhancements to make your life richer, in more ways than one. For instance, there are record and file locks, semaphores to help manage multiuser/multitasking data, and automatic disk recovery for better reliability.



Bill Gates, Chairman of Microsoft Corp.

Which makes XENIX 286 an ideal Operating System for applications software development. Over one hundred and fifty customized business software packages have been developed for XENIX, making it one of the most useable multiuser Operating Systems available today.

For Microsoft, the unique advantages of Intel's 80286 microprocessor were most appealing. As Bill Gates, chairman of Microsoft, said, "On-chip memory management and protection offered by the 286 ensures code compatability and makes it easy to port XENIX between different OEM systems."

In addition, its ability to run in fast 8086 mode makes the 286 the only processor that can support both XENIX 286 and MS-DOS without additional hardware. "With this ability, users get the best of both worlds in one piece of hardware," said Gates.

What's more, it's the only microprocessor architecture with the flexibility to offer you entry at whatever level of integration you choose: chips, boards, or complete systems (including, of course, our new 286/310 supermicro system.) So you get to market when you want, in the way you want.

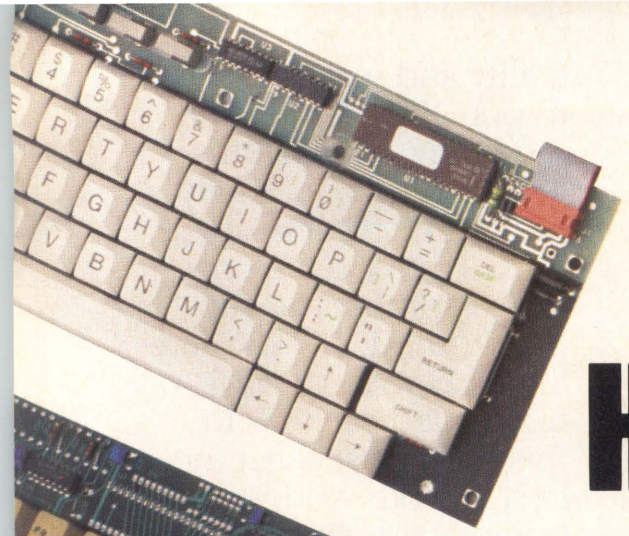
Why not dig a little deeper, and call us toll-free at (800) 538-1876. In California, (800) 672-1833. Or write Intel, Lit. Dept. #S9, 3065 Bowers Avenue, Santa Clara, CA 95051.

And don't be surprised if you hear shouts of "Eureka!" coming from your product development team.

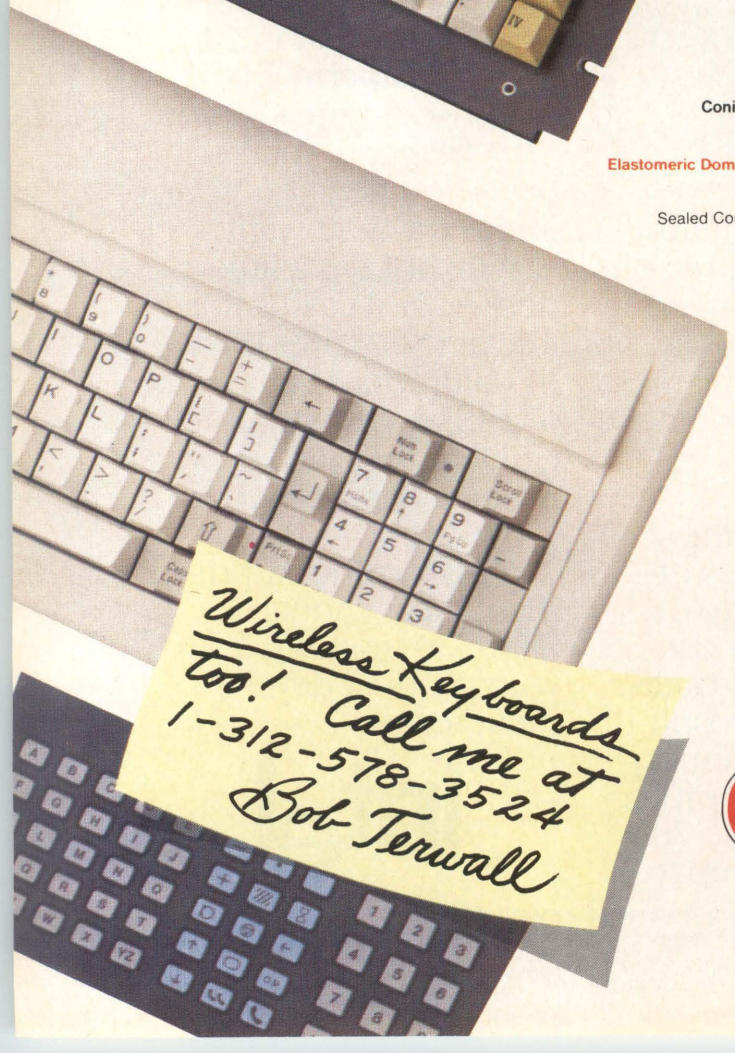
intel®

**XENIX and MS-DOS are trademarks of Microsoft Corporation. UNIX is a trademark of Bell Labs, Inc. AT&T is a trademark of American Telephone & Telegraph.* *© 1984 Intel Corporation.*

CIRCLE 13



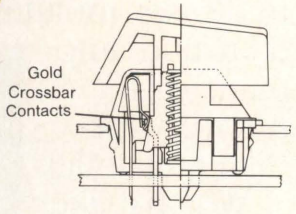




You choose from all the best technologies at

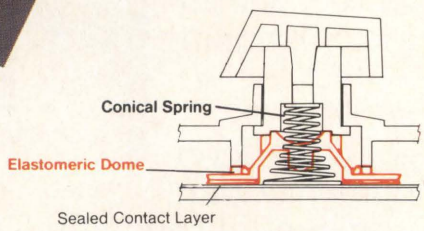
KEYBOARD HEADQUARTERS

DIN compatibility in the most advanced keyboard technologies.
Only from Cherry: Where we are ready today to meet 1985 ergonomic standards.
Only from Cherry: Where we make all the most cost-effective technologies for your application.
Only from Cherry: Your Keyboard Headquarters. Worldwide.



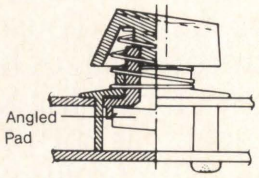
Full Travel Hard Contact

Gold crossbar contact configuration relied on in millions of applications worldwide. A first from Cherry more than a decade ago. Now proven and improved to meet the low-profile challenges of the 80's.



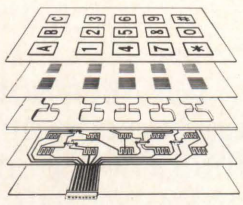
Full Travel Sealed Contact

NEW: A next generation keyboard that combines full travel with the quality and reliability of sealed silver contacts for long life, low cost. **Conical steel spring action** with linear feel... **or elastomeric action** with tactile feel.



Full Travel Capacitive

Pad capacitive in a uniquely simple design requiring only five parts and a snap-in angled pad. Another keyboard technology alternative from Cherry... the folks with the *most* logical alternatives.



Flat Panel Membrane

Unlimited design options thanks to our state-of-the-art production techniques and in-house fabrication. Thin, lightweight, reliable, low cost. All this and sealed contacts, too!

Wireless Keyboards too! Call me at 1-312-578-3524 Bob Terwall

SEND FOR COMPLETE DATA

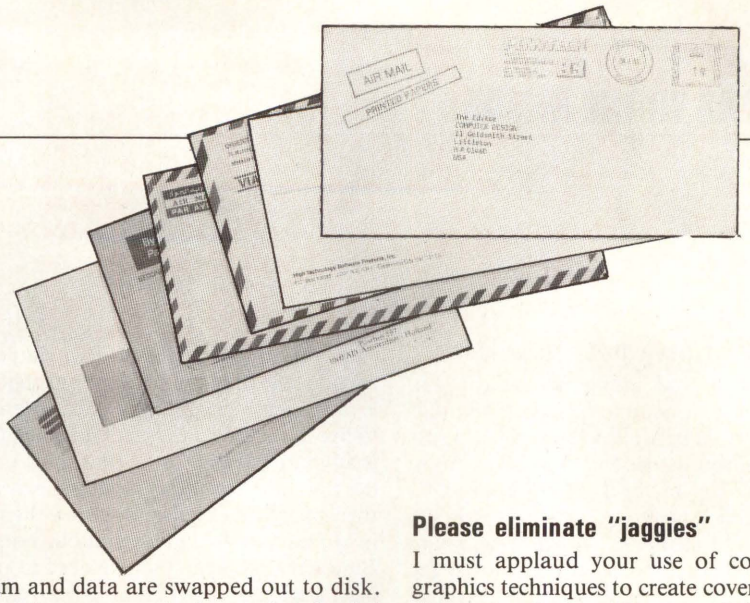
CHERRY®

KEYBOARD HEADQUARTERS

CHERRY ELECTRICAL PRODUCTS CORP.
3614 Sunset Avenue, Waukegan, IL 60087 • 1-312-578-3500

CIRCLE 14

LETTERS TO THE EDITOR



Unix: a closer look

I have enjoyed reading the April issue, especially the items concerning the Unix operating system. Sam Bassett's article, "Operating Systems Contend for Position as Industry Standard" (p 42), reflects a good grasp of the matter. His description of Unix as "endlessly modifiable" highlights a strength of its modular design. OEMs are using Unix as a meta-environment (an environment for building specific environments) to develop vertical market applications such as CAE systems, fault-tolerant systems, secure systems, business systems, and realtime SCADA systems. When desirable, it has been straightforward to add demand-paged virtual memory, record locking, and bit-mapped windowing.

Naturally, Bassett would be remiss if he did not point out the flaws in the Unix operating system. Its user interface was designed over 10 years ago for low baud rate teletypewriters and dumb CRTs. Until five years ago, it was lauded as a paradigm of user-friendliness. But the advent of microcomputers for the masses, with software designed for a whole new class of users, has changed that.

Hence, the article contains a few unfortunate statements concerning the Unix operating system, such as "nobody can be expected to understand [Unix] in a finite period of time." Because understanding is a qualitative abstraction, the same statement can be made of any operating system whose code will not fit on the back of an envelope.

In addition, "[Unix] is written in C—a far cry from Cobol, in which most applications are written," and "The whole system has a bad habit of changing completely every time a new system programmer works on it." First, microcomputer software applications written in Cobol are rare. They are almost exclusively written in Pascal, Basic, or assembler. Now, C is becoming the language of choice. Ask Digital Research, Microsoft, Lotus, or VisiCorp. Second, the system changes fault lies with system administrators, not Unix. Uncontrolled change is risky regardless of the operating system involved.

Bassett went on to say that "the supplied documentation, by its sheer weight and incomprehensibility, is guaranteed to give an English major a hernia and hysterics," and that "when [a program] must be halted for some reason, the entire program and data are swapped out to disk. When it can run again, it must be reloaded from disk and reconnected to the operating system."

In reality, large operating systems perforce have extensive documentation. Look at the number of manuals supplied with Multics or AOS. Technical manuals are notorious for misjudging their audience. The Unix System V manuals are a marked improvement over the past. And, now that AT&T has a desktop micro, the 3B2, I suspect we will see manuals that satisfy the needs of microcomputer users. In the interim, there are numerous books that will fill this gap.

Also, multitasking systems must either swap or limit the size and number of tasks that can be executed concurrently. Swapping definitely degrades performance in the face of low performance drives. But, the situation is not as dark as Bassett paints it. Unix System V swaps a waiting program only when there is insufficient memory to load the current program. Since text (instructions) is shared, this segment of program may not get swapped at all.

Fred Christiansen
Motorola Microsystems
2900 S Diablo Way
Tempe, AZ 85282

Don't ignore a good chip

I can understand Surendar Magar's desire to push TI's TMS320 in his article, "Signal Processing Chips Invite Design Comparisons" (Apr 1984, p 179). It is a very nice chip and I use it. However, I cannot understand how one can ignore, in an informative article, the existence of a high performance chip such as the Fujitsu MB8764, with such features as a 100-ns cycle, microprogramming, and a 256-word memory. This chip has been available for some time and can hardly be overlooked.

Gideon Keydar
Scitex Corp Ltd
PO Box 330
43103 Herzlia B, Israel

Please eliminate "jaggies"

I must applaud your use of computer graphics techniques to create cover art for *Computer Design*. However, I feel that you are giving computer graphics a bad name by continuing to allow your artists to use equipment that produces output with severe spatial aliasing. I am sure many of your readers are becoming disenchanted with computer graphics as an output medium because of the "jaggies" in your cover art.

There are many anti-aliasing techniques available and, undoubtedly, one of them would be appropriate for your system. I would encourage you to explore these techniques as they could significantly improve the quality of your covers without incurring excessive cost. Also, computer graphics neophytes will not get a negative impression before they have a chance to experience higher quality displays.

L. Jay Bass
Burroughs Corp
11010 Roselle St
San Diego, CA 92121

Other readers have questioned our use of relatively low quality graphics techniques when today's most advanced equipment allows much higher resolution. For this reason, beginning in September our covers will be created using a wider range of computer graphics techniques—including three-dimensional effects. One problem, of course, is that computer graphics can now create pictures that are almost indistinguishable from hand drawings or photographs of solid objects. Our dilemma, therefore, is to show textually that we are using computer graphics, while also availing ourselves of the full potential of the technology. For this reason, we may sometimes deliberately introduce "jaggies" into our cover art, even though the equipment is capable of eliminating them. Thus, we would be creating the high-tech equivalent of brush marks in an oil painting.

Michael Elphick
Editor in Chief

(continued on page 24)

(continued from page 23)

They don't have to be whiz kids

I read with great interest the editorial in the February issue (p 11) on Computer Literacy. [The following are excerpts from a short article Mr Eisenberg sent to members of his local school board.]

"The time has come the walrus said..." With this phrase Alice was swept into Wonderland. Today, we are all being swept into a different Wonderland—Microland. Computers will continue to pervade many facets of our lives. Our children are exposed to them before they enter grade school (remember the video game hooked to your living room TV).

In our grade schools, "computer literacy" is the new buzz phrase all educators are striving for. Everyone wants it—but how many know what *it* is? I have heard people at all levels of our educational system echo the theme "Computer literacy is necessary to survive in our technological society." Without it, they fear that our kids will be somehow left out. My question is, left out of what?

Are we forgetting the primary purpose, perhaps the only purpose, of our schools is to teach our children the skills necessary to live in the world they will see when they graduate? Now, I know that most of us do not understand computers; we feel threatened by them, and we see them as invading our workplace and our homes. But does that mean that we need to raise a generation of programming wizards? Or does it really mean we must teach our children to function in a world that will use computers much the way we use TV's today?

I have been working with computers for over 20 years; I have been using a TV for longer than that. I am fully willing to admit that I have no real understanding of how a TV works, but this lack of knowledge has not proved a significant handicap to watching what I want, when I want. The TV is a *tool*—one designed to provide entertainment.

The computer is also a tool, a multifunctional tool perhaps, but a tool nonetheless. For most of us, knowing how to use a tool is both necessary and sufficient. Understanding how it works from the inside out, or how to program it is unnecessary. Providing this functional understanding of how to use computers is what I mean by "computer literacy."

I think there is a problem in assuming that all children need to know how to program computers to survive in the world of 1988 to 1995. That kind of perception leads us into wanting to train all our kids in computer programming. What is the role of the computer in our schools? What functions can it serve? How can we provide the right form of "computer literacy"? And, perhaps most importantly, how can we use its power to augment the educational process? Clearly, turning our children into a generation of hack programmers is *not* the answer to these questions!

Just as Alice strayed into Wonderland and experienced both exciting and frightening adventures, so too can the novice wander in Microland with equal delight and profit. How good or how bad the experience is depends only on the quality of the planning done beforehand and the understanding of the real needs of the organization.

A. M. Eisenberg
E. I. DuPont de Nemours & Co
Wilmington, DE 19898

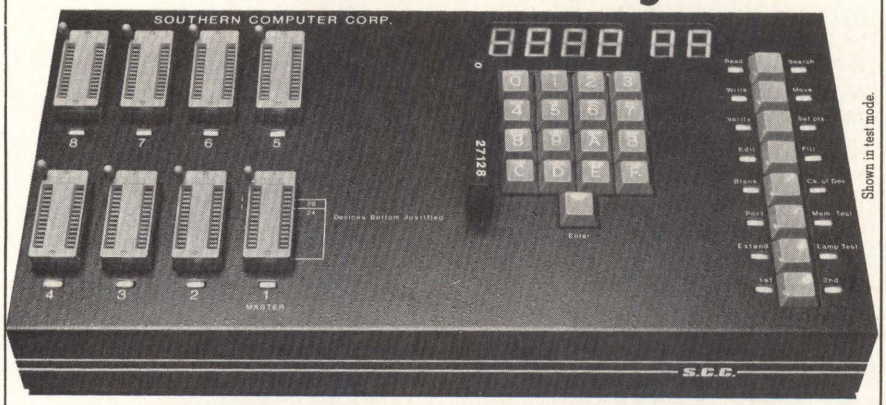
Another view on reliability

In response to Michael Hordeski's letter in the June 1 issue, my point is that reliable designs require a lot of careful work. There are many potential problems, and some of them are quite obscure. I do not think reliable microprocessor design is a lost cause—far from it. Many of them are very good.

The use of normal MTBF (MIL STD 217) reliability prediction simply takes into account the reliability of each individual component. This does not account for race conditions, noise problems, or even designs that don't work. The MTBF method of reliability prediction is simplistic and is only one small part of a truly reliable design. If we could only use this method as Mr Hordeski suggests, anyone could beat us.

Carl Oppenheimer
12 Depot Rd
Kingston, NH 03848

The Cost Efficient EPROM Programmer!



DISPLAY ☐ Bright 1" high display system ☐ Progress indicated during programming ☐ Error messages

KEYBOARD ☐ Full travel entry keys ☐ Auto repeat ☐ Illuminated function indicators

INTERFACE ☐ RS-232C for data transfer ☐ 110-19.2K baud ☐ X-on X-off control of serial data

FUNCTIONS ☐ Fast and standard programming algorithms ☐ Single key commands ☐ Search finds data strings up to 256 bytes long ☐ Electronic signatures for easy data error I.D. ☐ "FF" skipping for max programming speed ☐ User sets memory boundaries ☐ 15 commands including move, edit, fill, search, etc. functions ☐ Extended mode reads EPROM sets

GENERAL ☐ Stand alone operation, *external terminal not needed for full command set* ☐ Total support ☐ 28 pin sockets ☐ Faulty EPROMS indicated at socket ☐ Programs 1 to 128K devices ☐ Built in diagnostics ☐ No calibration required ☐ No personality modules to buy ☐ Complete with 128K buffer ☐ Only

$995.00 COMPLETE

Dealer inquiries welcome.

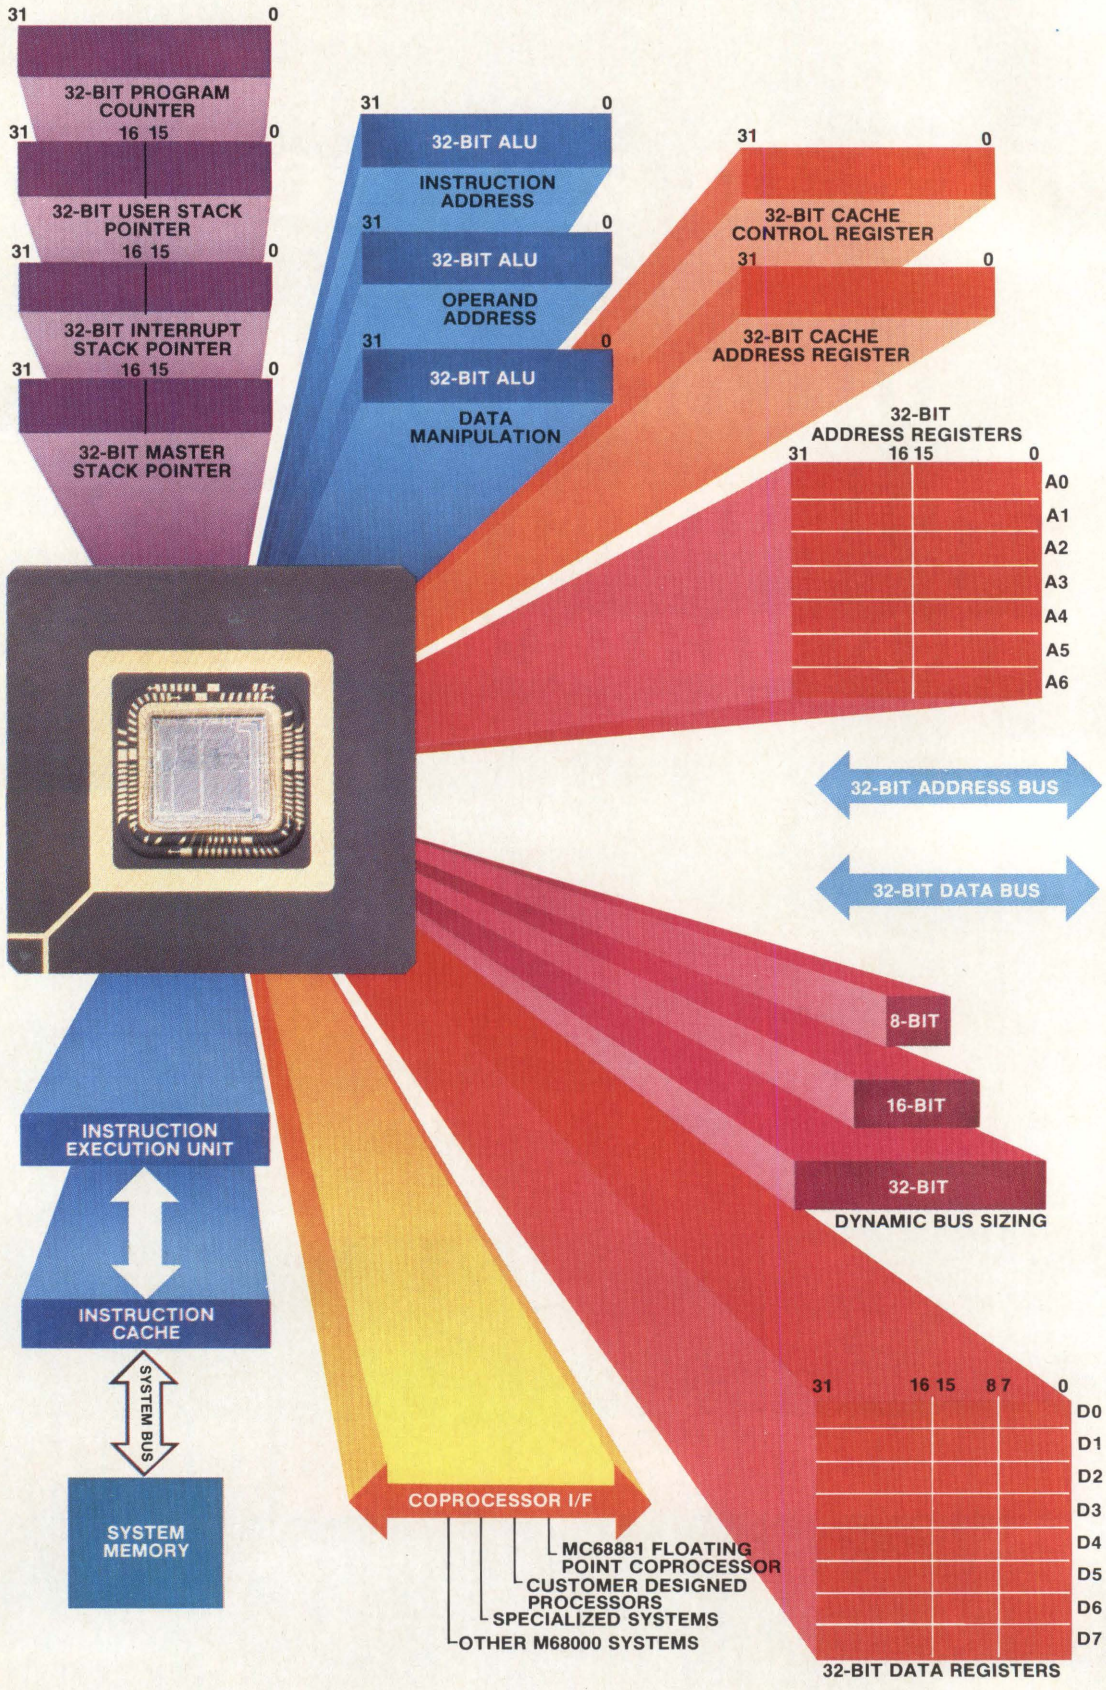
SOUTHERN COMPUTER CORPORATION
3720 N. Stratford Rd., Atlanta, GA 30342, 404-231-5363

Once in a generation:

introducing the 32-bit microprocessor performance standard.



Unleash the potential of your
the complete 32-bit



new system with the MC68020: microprocessor.

The new performance standard.

Motorola's new MC68020 performs at speeds typically 400% of the established standard of comparison, the MC68000. It's up to ten times faster in dedicated 32-bit applications.

No other 32-bit MPU makes this extensive a leap in performance improvement. At 16.67 MHz the MC68020 typically runs at 2.5 MIPS for integer processing. MIPS rates several times typical are achievable in dedicated 32-bit applications.

The advanced two-micron HCMOS manufacturing technology which allows this unparalleled performance also results in very low power dissipation. In fact, the MC68020 consumes less power in a system than the original MC68000.

The MC68020 creates opportunities you've never had before--opportunities to unleash the full potential in your 32-bit MPU-based systems because it sets the standard for 32-bit microprocessors. And, because it's the first complete 32-bit microprocessor available, more than just a 16-bit design on a data bus stretched to 32 bits. A detailed look at the architecture reveals this totality.

A fully compatible M68000 Family member.

Yes, the MC68020 has features new to the M68000 Family to maximize its true 32-bit capabilities.

Yes, it's an all new design built with advanced, highly manufacturable HCMOS technology.

And, yes, it's a fully-compatible member of the M68000 Family of MPUs and peripherals. All user object code written for previous M68000 Family MPUs executes without revision. In fact, MC68020 enhancements allow it to run more than three times faster.

Family compatibility is further enhanced by dynamic bus sizing, which supports the use of 8-, 16- and 32-bit ports in 68020-based systems. In fact, the MC68020 can be used in existing 8- or 16-bit systems.

New features enhance 32-bit architecture.

The MC68020 design is new, however its architecture is based on the proven M68000 Family 32-bit register set. And, the MC68020 is highly enhanced.

On-board instruction cache speeds operation and provides increased multiprocessing efficiency. The coprocessor interface allows direct expansion of the architecture off the MC68020 chip to coprocessors or customer-specified processing systems.

New addressing modes, new instructions and a 32-bit barrel shifter support new capabilities. Operating system efficiency is improved with a 32-bit program counter.

These enhancements and more optimize the MC68020 for 32-bit operations.

Design support brings projects together, fast.

Making the most of your new 32-bit design opportunities with the MC68020 is simple and effective with the backing of powerful new hardware and software support from Motorola.

The Benchmark 20™ evaluation system has been developed as a maximum environment testbed for resultant software. For initial software development, cross-support packages under both the UNIX™-derived System V/68™ and the real-time VERSAdos™ operating systems run on standard Motorola VME/10™ and EXORmacs® hosts.

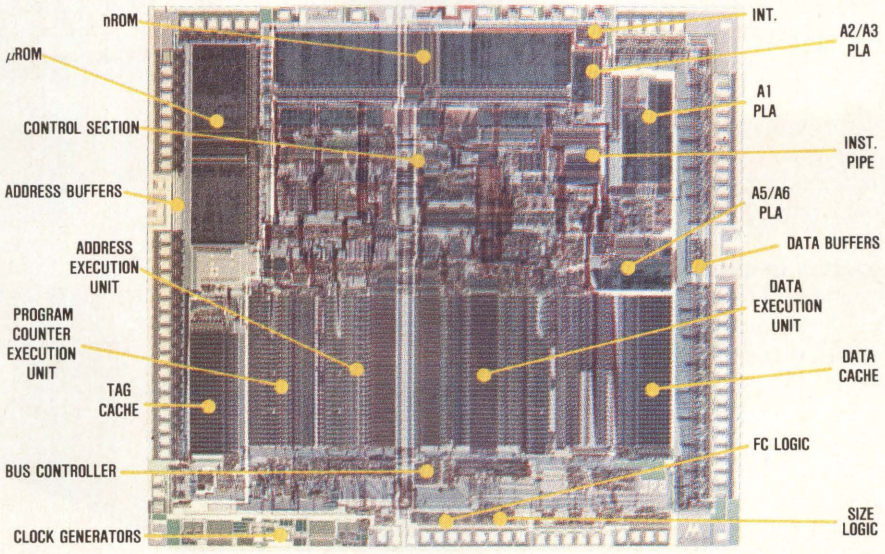
You'll find MC68020 designs a breeze with Motorola's advanced development tools--real time emulation and bus-state analysis with the HDS400 development system.

Move up to the MC68020.

The opportunity to design new-generation systems around the MC68020 and the M68000 Family is yours today. Marketplace attention will be focused directly on the growth-oriented companies that take advantage of this opportunity. Motorola's sales engineers and field applications specialists are available and equipped to assist you in moving up to the new 32-bit microprocessor performance standard. Contact one of them today.

Additional technical information is available by writing or sending the completed coupon to Motorola Semiconductor Products Inc., P.O. Box 20912, Phoenix, AZ 85036.

Benchmark 20, SYSTEM V/68, VME/10 and VERSAdos are trademarks of Motorola
EXORmacs is a registered trademark of Motorola
UNIX is a trademark of AT&T Bell Laboratories.



MOTOROLA



TO: Motorola Semiconductor Products Sector
P.O. Box 20912, Phoenix, AZ 85036

Please send me more information on MC68020.

193CD080084

Name_____

Title_____

Company_____

Address_____

_____

Call Me. Phone:_____

Because You need to run FORTRAN programs 10X to 100X faster...



Mini-MAP makes it practical to apply array processing to general-purpose scientific and engineering computing.

Practical in terms of use: Mini-MAP's compiler allows you to program the array processor directly in FORTRAN. An assembler, a linker, and a debugger are also part of the package. Plus you can use our library of over 250 highly optimized scientific subroutines.

Practical in terms of throughput: Because it is an array processor, Mini-MAP increases the computing speed of a mini or super-mini computer as much as 10 to 100 times. Where it takes a typical minicomputer minutes to perform tasks such as image rotation, Mini-MAP reduces interactive response times to seconds. Your computer may require hours to perform each step of a trial-and-error-process such as simulation, but Mini-MAP, can zip through in mere minutes.

DEC, PDP-11, LSI-11, and VAX-11 are trademarks of Digital Equipment Corp.
Mini-MAP is a trademark of CSPI.

We built Mini-MAP.™

Practical in terms of cost: Mini-MAP is available as an economical, four-board set or as a packaged system. Now, with Mini-MAP, OEMs can offer their customers a better product at lower costs. Mini-MAP's low power demands, small size, and high reliability make the package extremely attractive. And end users will find our FORTRAN compiler and other software tools minimize program development costs.

Some practical things to know about Mini-MAP: 32-bit DEC™ floating point format ☐ Interfaces to DEC PDP-11, LSI-11, and VAX-11 series ☐ Up to 16 MBytes of data memory ☐ 1024 x 1024 2-D FFT in 8.8 seconds ☐ Extensive software tools plus dedicated applications assistance including training, convenient parts depots, and field service staff support our worldwide installations.
To find out how Mini-MAP can work for you, **call toll free 1 800 325-3110.**

cspi
THE ARRAY PROCESSORS

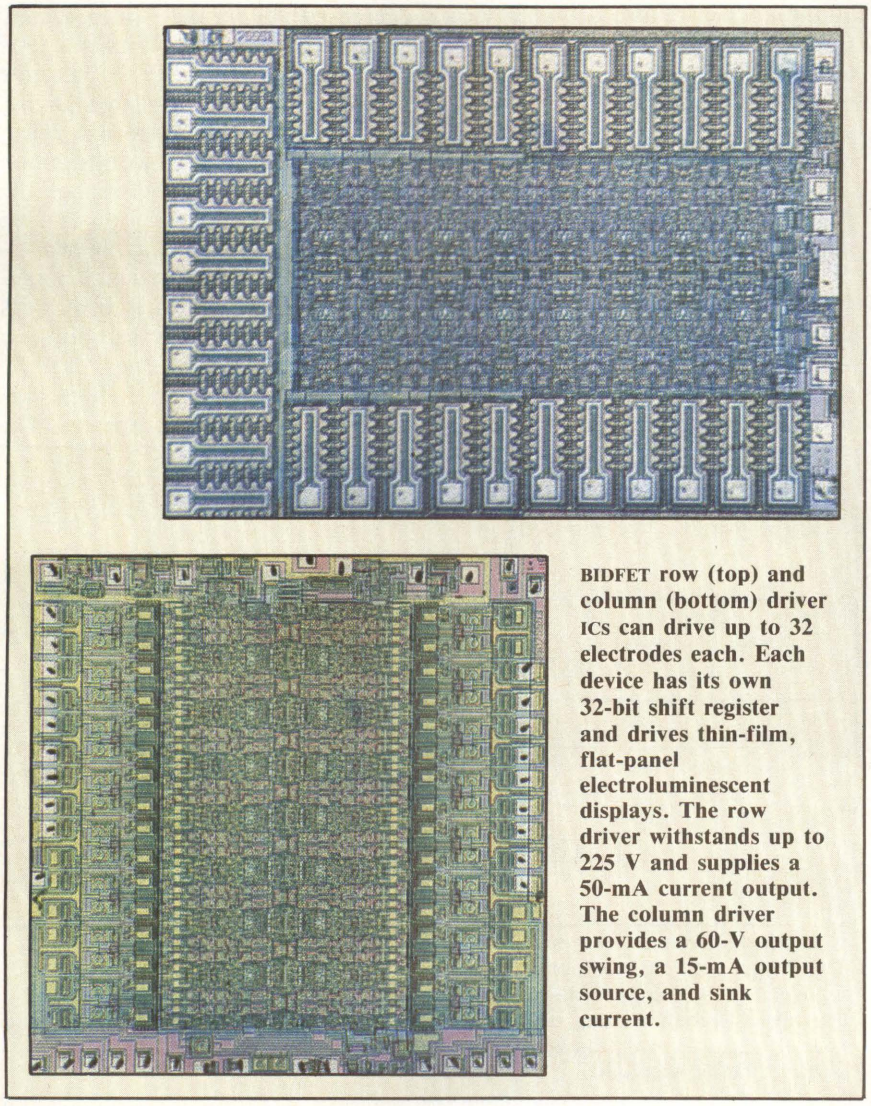
40 Linnell Circle, Billerica, Massachusetts 01821 • 617/272-6020 • TWX 710-347 0176

CIRCLE 17

Power and logic devices are merging on the same chip

A revolution is brewing in the field of power semiconductors. Sheer brawn is giving way to brawn and brain, as power semiconductor devices gain greater intelligence in the form of logic on the same chip holding the power transistor(s). Advances in merging bipolar and MOS technologies have made this possible. Although the technology of joining logic and power on the same chip is still very young, it holds immense implications for circuit and system designers. With decentralized microprocessor power already widely available, the next step is likely to be truly intelligent power devices with microprocessors on the same chip holding the power device, leading to distributed power control.

Much of the success of intelligent power IC technology rests on the ability of process engineers to merge such unlikely partners as CMOS devices, generally used for logic functions, with high speed and high power bipolar and double-diffused MOS (DMOS) devices. Until a few years ago, such merged technologies were difficult to master and uneconomical to use. However, with the advent of improved lithographic etching and ion-implantation equipment, as well as greater experience in making merged processes work, intelligent power ICs made with mixed technologies are looking economically more attractive.



BIDFET row (top) and column (bottom) driver ICs can drive up to 32 electrodes each. Each device has its own 32-bit shift register and drives thin-film, flat-panel electroluminescent displays. The row driver withstands up to 225 V and supplies a 50-mA current output. The column driver provides a 60-V output swing, a 15-mA output source, and sink current.

Designers are discovering that not only does the merger of logic and power technologies on one chip produce smaller size ICs than if two separate devices are used, it also means lower overall costs, increased performance levels, and higher reliability. Moreover, circuit designs are simplified considerably.

The ultimate goal is to create intelligent power devices with the fastest possible switching speeds and the most efficient operation, at the lowest possible cost. This pursuit of theoretical performance limits has MOSFET power device designers chasing after all sorts of cell layout geometries, to get the most performance from the least amount of silicon area. The goal is to achieve maximum packing densities and the lowest MOSFET on resistance. The larger the ratio of the FET's source width to area, the more active silicon in the form of individual cells that can be achieved on a given chip.

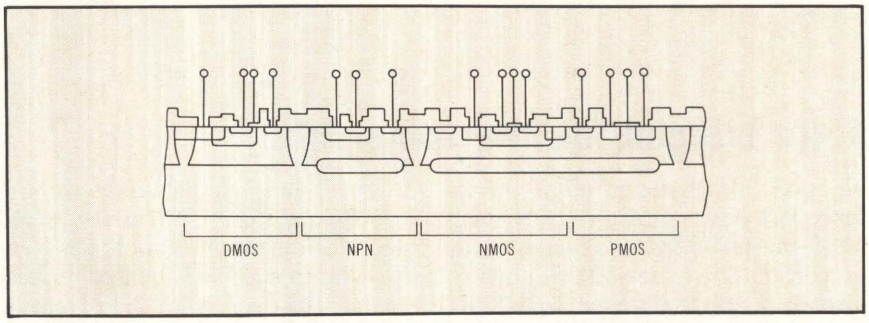
The shape of things to come

Thus, it is not surprising to see a plethora of MOSFET cell shapes emerging, with square and hexagonal cells predominating. The square cell is the choice of such device manufacturers as Motorola Semiconductor Products Inc, Siemens Inc, and Siliconix Inc. Japanese manufacturers like Hitachi Ltd, Nippon Electric Corp, Toshiba, and Matsushita are also opting for the square-cell

(continued on page 30)

Power and logic devices

(continued from page 29)



The BIDFET process brings together on one chip double-diffused MOS (DMOS), npn, NMOS, and PMOS transistors. It allows the merging of precision-control and self-isolated CMOS logic devices with high voltage interface circuitry through the use of standard junction-isolation techniques.

approach. International Rectifier Corp, on the other hand, feels that the hexagonal cell it pioneered is superior; so does General Electric Co, Unitrode Inc, and RCA Semiconductor Inc. Others, like Texet Corp, are using triangular cells, while Supertex Inc uses an interdigitated approach. Besides the square cell, Siliconix also uses the interdigitated approach.

Regardless of what cell shape is used, all are achieving record levels of silicon-area utilization, record high voltage capabilities, and new lows in MOSFET on resistance. And, they are doing it while successfully mixing with other process technologies. The Texas Instruments, Inc (PO Box 225012, Dallas, TX 75265) BIDFET driver ICs for electroluminescent flat-panel displays are but one example.

The SN7551/7552 row drivers and SN7553/7554 column drivers are made by the BIDFET process that mixes bipolar, DMOS, and FET transistors on the same chip, hence its BIDFET name. The BIDFET process merges precision-control, self-isolated CMOS logic with high voltage interface circuitry on a common monolithic substrate, by using standard junction-isolation techniques. It solves the high voltage limitation problems of conventional ICs while retaining their logic capabilities. BIDFET devices have been produced with working voltages of up to 250 V and breakdowns of more than 300 V. This is achieved by replacing the conventional bipolar output stage with a DMOS transistor.

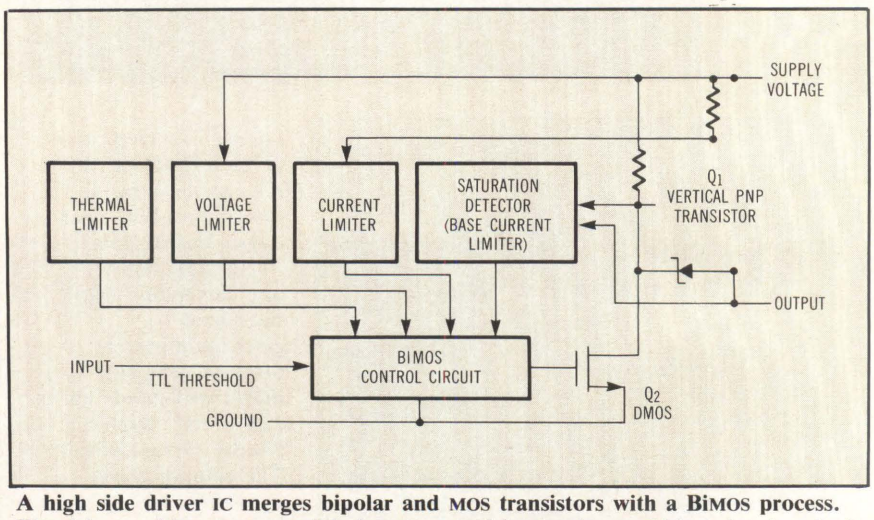
Driving 32 electrodes, the row drivers can withstand up to 225 V and provide a 50-mA output current. Input is CMOS compatible. The column driver IC also drives 32 electrodes with a 60-V output swing, and a 15-mA output source and sink current. It features high speed serially shifted data, and totem-pole latched outputs.

The junction-isolated process TI uses in making the BIDFET drivers is gaining in popularity over another process known as self isolation. The former process provides higher onchip current (about 10 to 20 vs 2 to 4 A), and lower on resistances. The trade-off is a lower maximum voltage rating (about 150 vs 400 V), and an additional processing step for an epitaxial layer.

Yet another mixed process yielding intelligent power devices is the Bi-MOS process used by Motorola. This process not only allows power devices with minimum power dissipation, high efficiency, and direct interfacing with external digital circuits, but also the ability to block both negative and positive high voltage transients.

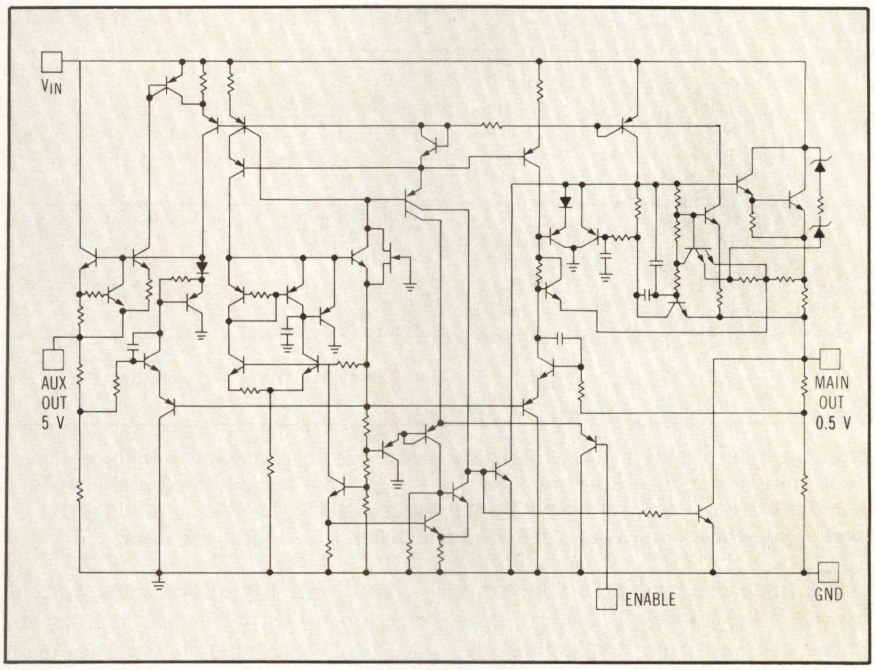
In this power process, a vertical pnp bipolar transistor is used at the output of a power device. As a result, die area is minimized through the presence of a substrate collector that eliminates both the top collector contact area and the lateral spacing needed to maintain high breakdown voltages. Collector series resistance is decreased by using the package header as the output. With this type of structure, better gain is possible than with lateral transistors having the same die area. There is a price, however. Latching is more common due to the forward biasing of the epitaxial-substrate junction during device operation, leaving the transistor's output collector at a high potential. This can be alleviated using MOS transistors that are less susceptible to latching than bipolar ones.

Motorola has made a high side driver with the BiMOS process. The



A high side driver IC merges bipolar and MOS transistors with a BiMOS process. The driver, with TTL-compatible inputs, can drive loads to within a few hundred millivolts of the supply voltage, and features voltage and thermal shutdown capability, current limiting, and a saturation detector. Lateral pnp transistors and laterally diffused MOS transistors within the device block negative, as well as positive, 125-V transients.

circuit has TTL-compatible inputs and can drive loads to within a few hundred millivolts of the supply voltage. The device features thermal and voltage shutdown capability, current limiting, and a saturation detector. Lateral pnp bipolar transistors and laterally diffused MOS transistors block negative and positive transients of up to 125 V.



One of the most intelligent monolithic voltage regulators is the LT1005. This dual-output IC has two independent 5-V regulators and is logic controlled. The main regulator's output can be logically controlled by a TTL- or CMOS-compatible enable signal, with its 5-V output switchable to near zero by a logic LOW signal. Meanwhile, the other regulator works without being affected by the logic control signal or by any fault condition at the main regulator.

Intelligent regulators too

Motorola Semiconductor Products, Inc (5005 E McDowell Rd, Phoenix, AZ 85008) makes an entire family of intelligent power devices appropriately named SMARTpower. One such device is the MPC100, a 5-V fixed regulator with an 80-W output transistor. The device can deliver up to 10 A, has internal thermal protection and short-circuit protection, overcurrent limiting, and, thanks to the regulator's low saturation voltage, features a low differential-voltage operation of just 1.5 V at the maximum output current of 10 A. The device is intelligent enough that it becomes a variable voltage regulator when a pair of external resistors are added to it.

Yet another SMARTpower device is the MPC2005, an overvoltage and temperature-protection circuit made of a combined CMOS and TMOS process. (TMOS is so named because drain current is split into a T formation.) The device, which combines the best of CMOS and bipolar transistor characteristics, discharges up to 150-A peak currents and operates at up to 15 A of continuous anode current. It trips when voltage exceeds 6.2 V or when the junction temperature rises over 125 °C, to remove the voltage supply from the system to which it is connected.

One of the most intelligent monolithic voltage regulators is the LT1005 from Linear Technology Corp (1360 McCarthy Blvd, Milpitas, CA 95035). The dual output device has two independent 5-V regulators and is logic controlled by a TTL- or CMOS-compatible enable signal. Its 5-V output can be switched to near zero by a logic low signal. The auxiliary output is unaffected by the logic control signal, as well any fault condition at the main regulator circuit, and will continue to deliver its 5-V output. Such a dual structure allows a system to be designed such that the main 5-V supply can be shut down for power savings or protected in the event of failure, while the auxiliary supply continues to apply power to other critical circuits.

Fault protection is what distinguishes the IM2925 monolithic regulator from National Semiconductor Corp (2900 Semiconductor Dr, Santa Clara, CA 95050). The 5-V, 750-mA device with a small I/O differential of just 600 mA, automatically shuts down in case of a short circuit, thermal overload, or the presence of input transients anywhere from −50 to 60 V. Not only that, but it automatically resets itself and restarts after a delay that can be set by the user. That delay can range from milliseconds to seconds, depending on the value of an external capacitor that is connected to one of the IC's pins. The regulator continues operating with an input as high as 31 V. When a fault occurs, the regulator's pin 5 shows a normally high reset flag, which can be sensed by a microprocessor for switching over to another voltage regulator.

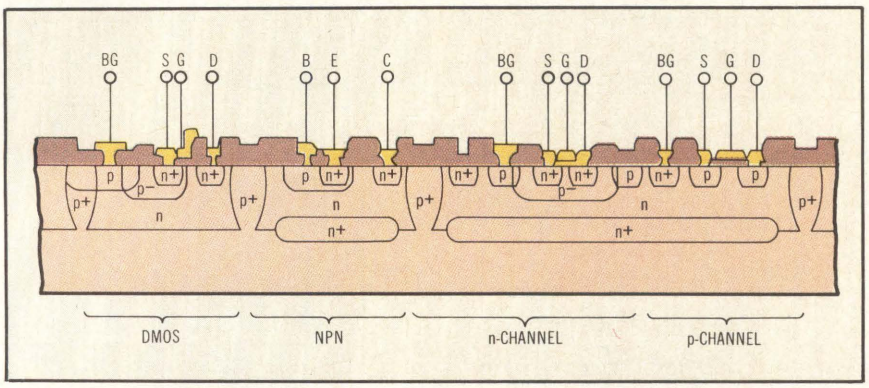
TI has used its BIDFET process to make an intelligent and adjustable monolithic voltage regulator supplying up to 1 A over a range of 5.25 to 200 V all in a TO-220 package. Unlike the firm's display drivers, most of the regulator's transistors are bipolar, and the output stage is a lateral DMOS transistor.

Intelligence also lends a hand in making monolithic voltage regulators more versatile, as can be seen in the Unitrode Corp (580 Pleasant St, Watertown, MA 02172) UC1834 three-terminal regulator. It includes many features that would normally require external circuitry for other three-terminal regulator ICs. For example, besides its high gain error amplifier and output-driver stage, the unit has onchip provisions for current sensing, remote-voltage sensing, and thermal protection. The regulator's output

(continued on page 32)

Power and logic devices

(continued from page 31)



CMOS logic devices (right) may share the same substrate space with high current DMOS transistors (left). Siliconix has used this process to build a controller IC with exclusive-OR logic to switch 12-A peak currents at 120 V and at 200-kHz rates. The process can handle 150-V potentials and 20-A peak currents.

transistor can serve as a sink or a source, allowing one device to serve as either a positive or a negative output supply. The chip furnishes 200 mA over ambient temperatures from −55 to 125 C.

Smarter power transistors

Thanks to a BiMOS process that General Electric Corp, Semiconductor Division (W Genesee St, Auburn, NY 13021) perfected, 500-V power transistors have been teamed with 20-V CMOS logic devices, on a single junction-isolated IC. There is only one drawback to such devices—the emitters of the bipolar transistors and the sources of the MOSFETs must be connected to ground for operation. This slightly limits the devices' flexibility.

GE achieved the 500-V breakdown rating by a technique known as lateral charge control. The technique permits lateral high voltage devices to be built in 7-μm, lightly doped n^- epitaxial layers grown on a p^- substrate. Unlike standard processes where the voltage breakdown of a transistor is a function of the thickness of the epitaxial layer, the breakdown voltage for the GE BiMOS devices is a function of the length of the charge-control drift region and the depth of the implanted impurity ions. The longer the drift region, the less the charge per unit area for a given voltage.

Until recently, if designers wanted to control power MOS transistors with 5-V logic signals, they not only needed special level shifting circuits, but also an additional supply voltage to drive those circuits. This is because it takes at least 10 V between a power MOSFET's gate and source terminals to get it into full saturation. The logic-level FET (LLFET) family of MOSFETs from RCA Corp, Solid State Division (Rte 202, Somerville, NJ 08876) has done away with all this. They can be driven into saturation with just 5 V between their source and gate terminals, sacrificing no speed, resistance or current-handling capability.

Teledyne Semiconductor Inc (1300 Terra Bella Ave, Mountain View, CA 94043) also makes level-shifting devices. Three dual power CMOSFET drivers translate a low level TTL or CMOS input signal into an ouput voltage swing within 25 mV of the power supply rails. The TCS426 has two inverting drivers, the TCS427 is a noninverting driver, and the TCS428 has one inverting and one noninverting driver. Without external capacitors to speed their operation, each device can swing a 1000-pF load to 18 V in just 30 ns. This combination of high speed and a wide voltage swing ensures that power MOSFETs are fully turned on or off, minimizing conduction power dissipation. At low supply voltages and with high threshold power devices, the rail-to-rail swing is particularly useful in minimizing power loss in switch-mode power supplies, motor-control systems, and dc-to-dc converters.

Because the gate-drive requirements of power MOSFETs have been reduced, the FETs are becoming attractive candidates as output devices for smart power ICs. For example, International Rectifier Corp, Semiconductor Division (233 Kansas St, El Segundo, CA 90245) takes a LED, illuminates a pile of photovoltaic cells with the LED's output, and triggers the gate of a MOSFET with the resultant electrical output of the photovoltaic cells. All of these functions are on one chip, and allow switching up to 1000 V. The device includes a bidirectional MOSFET known as bidirectional output-switch FET (BOSFET) and control logic on one chip.

The company also uses the photovoltaic principle in its solid-state switch with zero crossing (S^3X) solid-state relay chip. The IC contains an optical receiver, signal conditioner, zero-crossing detector, and a high power thyristor, all on one substrate.

Another example of a smart power switch is the experimental XPC1500 logic-to-power switch for Motorola, for switching up to 16 A, on command signals received from CMOS logic or TTL inputs. The firm is also developing the XPC1600, a pulse-width-modulated power IC that can handle 150-V potentials at a maximum current drain of 10 A. The MOSFET device is designed for switching-regulator and motor-control applications.

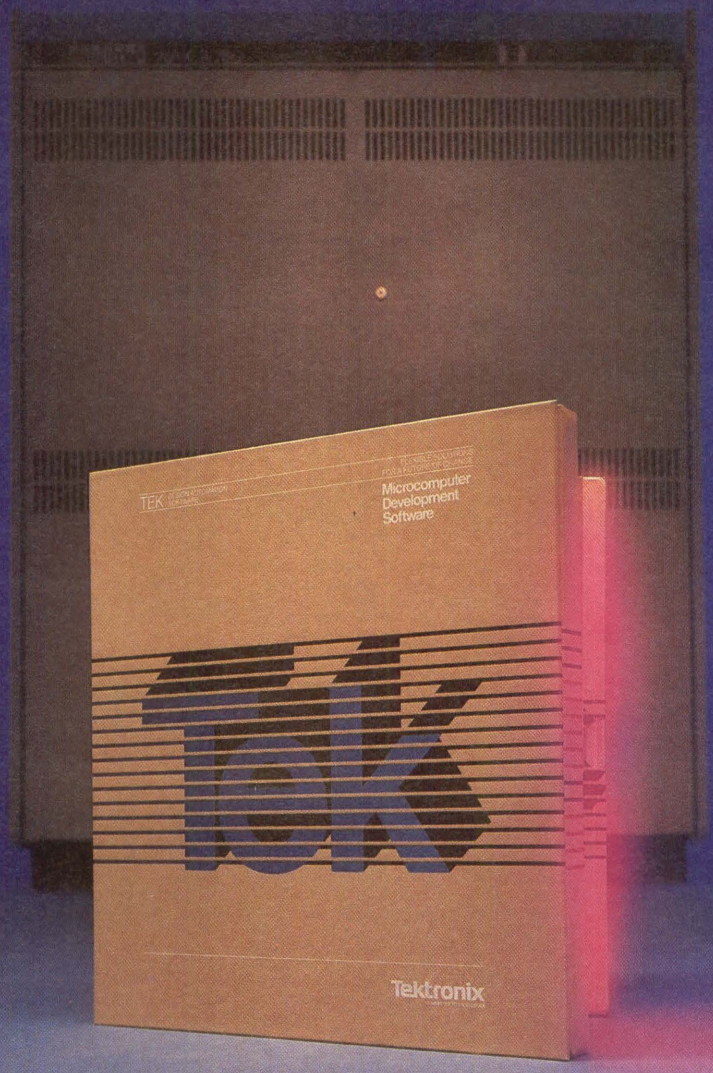
Siliconix Inc (2201 Laurelwood Rd, Santa Clara, CA 95054) has developed a controller chip to handle up to 150-V potentials and peak currents as high as 20 A. This chip combines DMOS and CMOS devices on one substrate for greater intelligence in power devices. The junction-isolation process employed uses n^+ channel power devices implanted in an n^- epitaxial layer, and CMOS devices built in a p

(continued on page 35)

TEK MICROPROCESSOR DESIGN SOFTWARE

FLEXIBLE SOLUTIONS FOR A FUTURE OF CHANGE

Tek software sheds new light on your VAX.



Turn your VAX™ computer into a powerful microprocessor development system.

Tektronix software. The same powerful tools that set the standard for high-level programming on Tek's 8500 series of microprocessor development systems are now available for use on your VAX. Get the sophistication of Pascal and C Language Development Systems (LANDS). Plus real-time emulation and debug when you integrate Tek's 8540 emulation unit to your system. Add 4105 Color Graphics terminals to access Colorkey+, Tek's single-key interactive user interface. All fully integrated with VAX-specific communications software.

Support from the first line of source code to the last line of debug.

Call your Tektronix sales engineer. With your VAX, we'll help you create a system that suits your engineering environment, and show you how Tek and your VAX have met the challenge of microprocessor software design.

Tektronix Microprocessor Development Systems.

Look at us now!

Call 1-800-547-1512

™ VAX is a registered trademark of Digital Equipment Corporation.

U.S.A., Asia, Australia, Central & South America, Japan: Tektronix, Inc., P.O. Box 1700, Beaverton, OR 97075. For additional literature, or the address and phone number of the Tektronix Sales Office nearest you, contact: Phone: (800) 547-1512, Oregon only: (800) 452-1877, TWX: (910) 467-8708, TLX: 151754, Cable: TEKWSGT.

Europe, Africa, Middle East: Tektronix Europe B.V. European Headquarters, Postbox 827, 1180 AV Amstelveen, The Netherlands. Phone: (20) 471146, Telex: 18312-18328

Canada: Tektronix Canada Inc., P.O. Box 6500, Barrie, Ontario L4M 4V3. Phone: (705) 737-2700

Tektronix®
COMMITTED TO EXCELLENCE

Copyright ©1984, Tektronix, Inc. All rights reserved. MIA-131

Circle 18 for Literature
Circle 19 for Sales Contact

It's easy to see why our new CRT series is an improvement down the line.

Improvements that give our HD-100 Series the good-looking features you expect from a quality CRT. You get sharper, cleaner lines and characters, plus improved focus and stability—even in the corners. And you can display over one million pixels on 15- or 17-inch screens, vertically or horizontally. In your choice of standard video or inverted displays, with line rates from 26 to 36 kHz. In effect, you can have enhanced visibility any way you want it, no matter whether your customers are processing words or graphics.

Our CRT line includes a number of other significant improvements. The HD-100 Series displays are 33% lighter than our previous HD Series. That adds up to easier handling, integrating and lower freight costs.

We also built in an efficient switch mode power supply that reduces power consumption for your customers while enabling a stable display over a wide voltage input range. In addition, our new HD-100 Series features an advanced CRT/yoke combination for better focusing. All our improvements are packaged neatly in a compact chassis design that offers convection cooling and makes service a breeze, and naturally, our CRTs are UL-478 approved and CSA certified.

All these enhancements put the focus on quality inside and out. High quality manufactured with great efficiency to give you a new, low price that's an improvement in itself. Call your nearest Ball sales office and ask to see our new HD-100 Series. The CRTs designed to give you a clear advantage with every line.

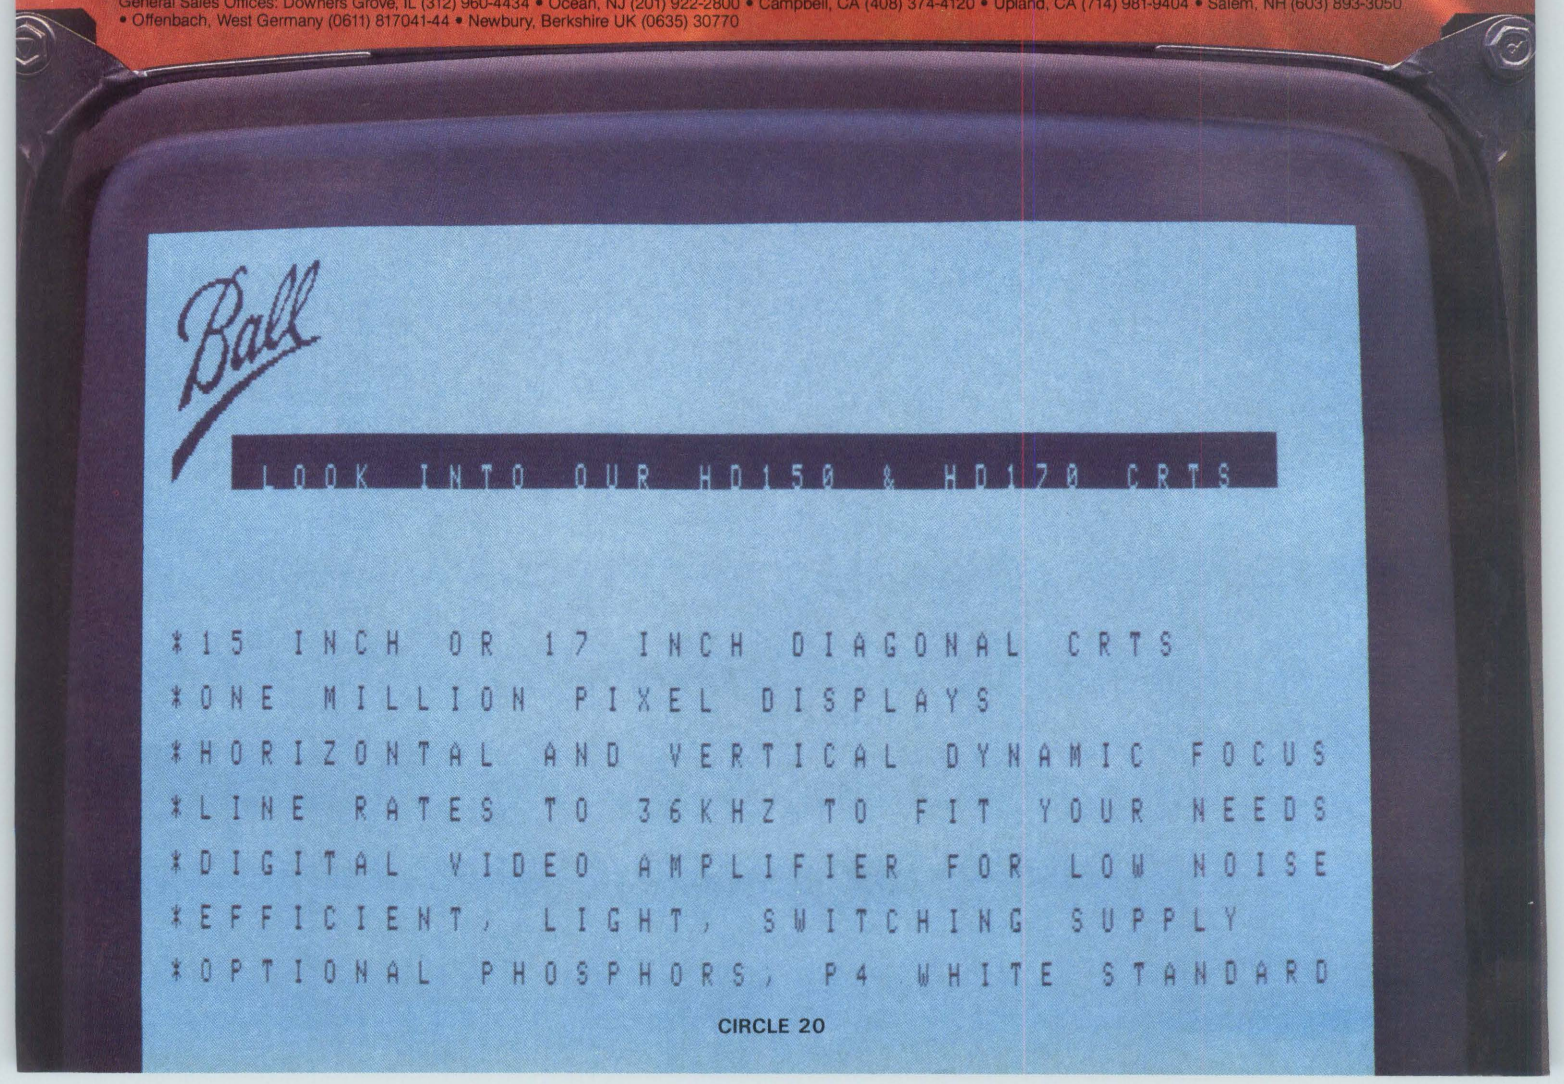


Ball Electronic Systems Division

P.O. Box 64376 • St. Paul, MN 55164

Bright ideas for better visibility.

General Sales Offices: Downers Grove, IL (312) 960-4434 • Ocean, NJ (201) 922-2800 • Campbell, CA (408) 374-4120 • Upland, CA (714) 981-9404 • Salem, NH (603) 893-3050 • Offenbach, West Germany (0611) 817041-44 • Newbury, Berkshire UK (0635) 30770



CIRCLE 20

Power and logic devices

(continued from page 32)

well. Power and logic transistors are isolated from each other by a p^+ region. The controller, which has a pair of power transistors with exclusive-OR logic on the same chip, can switch peak currents of 12 A at 120 V (per transistor) at 200-kHz rates. The company has also built a high speed MOS controller IC with blazing speeds. The device has a 100-ns dead time, which is about one-fifth that of previous devices, to allow it to operate up to 600 kHz.

One intelligent controller that simplifies the design of switching power supplies is the UC1846 from Unitrode. This pulse-width-modulated IC operates in the current mode to improve not only regulation and transient response, but can also be connected in parallel for high current applications.

Going the gate-array route

Semicustom and gate-array ICs are also being used to make intelligent power devices. This approach speeds the design turnaround time and product availability, and makes design prototyping simpler and less expensive.

With its TMG5002 gate array, Telmos Inc (740 Kifer Rd, Sunnyvale, CA 94086) has the highest gate-array voltage rating at 200 V. The complex array chip has 16 output stages, each of which is rated at 200 V and 40 mA, and a complementary output pair, also rated at 200 V and 40 mA. The TMG5002 has an assortment of 5-V CMOS transistors, 160 uncommitted gates, 23 D-type, dual-latch, flip-flops, six MOS or TTL I/O buffers, 10 input-only buffers, and four low voltage utility amplifiers.

The array's 200-V performance stems from the use of DMOS transistors and dielectric isolation processing. Not only does the entire CMOS section of the array reside in its own dielectrically isolated island, but each 200-V DMOS transistor (either p- or n-channel) does too. As a result, the transistors interface with 5-V CMOS logic as if they were discrete devices, forming, for example, a 200-V push-pull, three-state output.

Making it easier for users to try out the Telmos chip is a kit of parts that contains the three-state circuit that is driven by 5-V logic, as well as an n-channel and a p-channel 200-V 10-mA DMOS transistor. Each of the transistors has resistance of less than 600 Ω at 10 mA.

American Microsystems Inc (3800 Homestead Rd, Santa Clara, CA 95051) also has CMOS standard cells that allow combining of intelligent power circuits. The transistors have complementary output buffers that swing between 0 and 100 V and are of the parasitic substrate npn variety.

For really high power in an intelligent device, designers can turn to the Tuff chips from Sprague Electric Co's Semiconductor Division (115 NE Cutoff, Worcester, MA 01606). The ULN-2350C and ULN-2351C are rated at 80 V each, with the former and larger device handling up to 1 A, thanks to four npn transistors on its chip, connected in parallel. The oxide-isolated polysilicon resistors on these chips are dielectrically isolated from the remainder of the chips' circuits to accept up to 500-V potentials. As a result, the use of an external resistor and zener diode, hooked to one of the IC's pins, allows protection against 500-V transients.

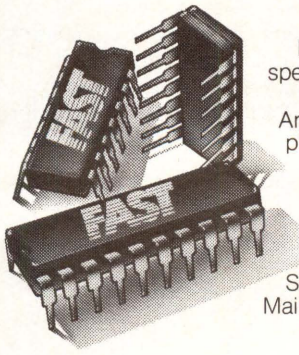
—Richard Parker, Contributing Editor

SYSTEM TECHNOLOGY

(continued on page 38)

September Preview
Watch for a special article on peripheral ICs

FAST Success.



In 1978 we introduced the FAST™ family of high-speed logic.
In 1983 we had our first million-unit customer.
And in 1984 we'll introduce FAST LSI into the marketplace, for a total of 100 FAST devices available.
For the FAST in your future, see your nearest sales office or distributor or contact the Product Marketing Dept., Fairchild Digital Products Division, 333 Western Avenue, South Portland, Maine 04106.

FAIRCHILD
A Schlumberger Company

FAST (Fairchild Advanced Schottky TTL) is a trademark of Fairchild Camera and Instrument Corporation.

Put an end to interference.

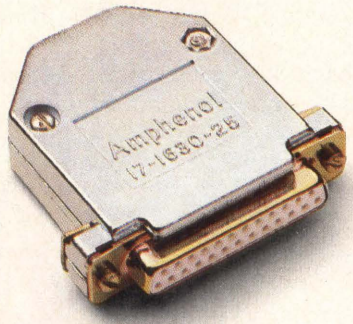
Interconnections are key points of susceptibility to EMI/RFI problems. In your battle to knock electronic noise out of your system, Amphenol Products may be your best ally. We've developed cost-effective anti-pollution solutions for your connection needs.

Using shielding to control radiated emissions, filtering for line conducted interference, and fibre optic technology for EMI/RFI protection, we can help you get your data straight. Call Amphenol Products, the new connector company, for more information, technical assistance and prices.



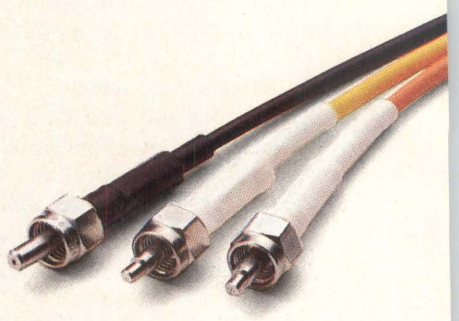
Shielded data communication cable assemblies. For compatibility with systems environments, Amphenol® interface bus cables are multi-shielded designs featuring double-shielded cable with additional shielding around the inner data bus lines. IEEE488-1975 style bus cables are terminated with 24-contact connectors. RS-232C/RS-449 style EMI/RFI shielded modem/data cables use tin-plated D-subminiature connectors with grounding indentations.

Reader Service Number 140



Shielded D-subminiature connectors. As the most commonly used input/output connector in business equipment, shielded D-subminiatures are a key element in FCC compliance. We provide radiation protection using a variety of backshell configurations, from conductive plated plastic, to stainless steel, to die cast alloy. Our backshells provide shielding coverage for all Amphenol®, Bendix®, and Spectra-Strip® RS-232C/RS-449 D-subminiatures.

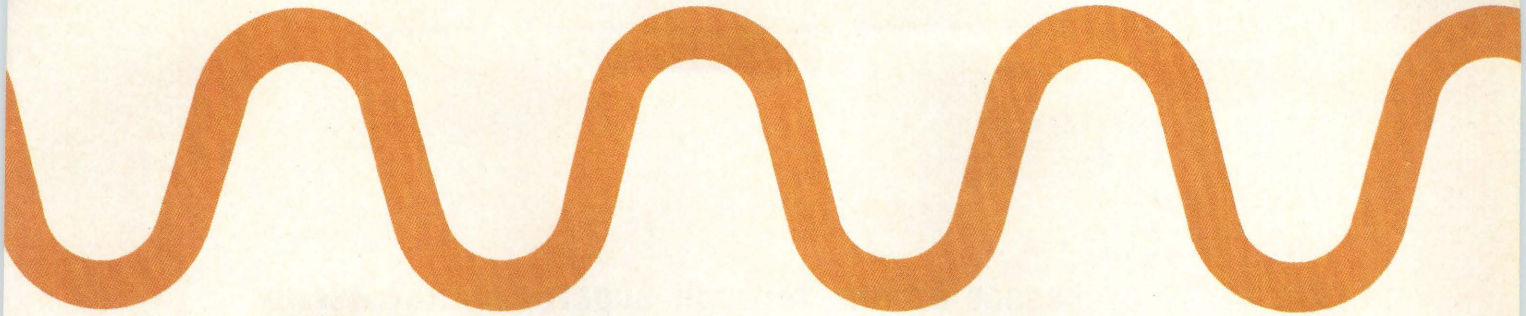
Reader Service Number 141

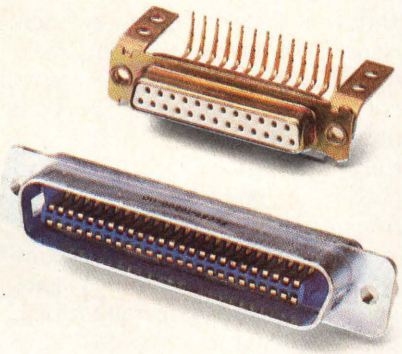


Fibre Optic Connectors. For the ultimate in reliable data transmission, we offer a broad range of fibre optic interconnection technologies and styles. Amphenol® and Bendix® connectors. From epoxy/polish and epoxyless/polish to epoxyless/no-polish. From single-channel, to duplex, to 29 channels in a single connector. From low-cost plastic connectors, to SMA styles. From rectangular styles, to environmental circular connectors. Also, for mounting semi-conductor devices.

Reader Service Number 142

©1984 Allied Corp.





Low-cost filtered input/output connectors. Our new cost-effective design uses stress-isolated NPO dielectric filter chips to eliminate conducted interference in each line. Available in a full range of D-subminiature and ribbon connector types and sizes, the connectors are intermateable with most existing industry standard non-filtered products. A pricing break-through, these filtered connectors can be matched to your insertion loss and band-pass requirements by capacitance value selection.

Reader Service Number 143



Filtered Military/Aerospace Connectors. Intermateable and intermountable with standard Mil-Spec circulars as well as MIL-C-24308 rectangulars. Amphenol® and Bendix® filtered connectors pass desired low frequency signals while attenuating line conducted EMI/RFI. These connectors are also designed for maximum protection from radiated EMI/RFI. Available in a broad range of filtering characteristics, shell sizes, styles and contact termination options.

Reader Service Number 144



Shielded planar products. Planar cables and mass terminated insulation displacement connectors are increasingly important in data transmission applications. Spectra-Strip® shielded planar cables, connectors and cable assemblies are your best choice for putting an end to garbled data. Cable is available in coax and 28 AWG versions, with integral aluminum/mylar shielding and optional drain wires. Stainless steel backshells provide total coverage and a low impedance path to ground for D-sub connectors.

Reader Service Number 145

Call the new connector company

1-800-323-7299

Amphenol Products world headquarters: Oak Brook, IL 60521

ALLIED Amphenol Products

Matrix array processor breaks through supercomputer barrier

Embracing both loosely and tightly coupled hardware architectures, the FPS-164/MAX matrix array coprocessor boasts peak performance from 33 to 341 million floating point operations per second. As such, its coupling with DEC VAX-11, IBM 3081, or Apollo Domain processors rivals performance previously attained by such supercomputers as the CRAY-1 or Control Data Corp's CYBER 180.

In addition to all this power, the system is easy to use. Users do not need extensive knowledge of its internal architecture to obtain optimum performance. A Fortran subroutine library provides the necessary hooks for application programmers to access both the matrix arithmetic elements as well as the floating point arithmetic unit. Assembly language facilities are also provided for those problems that require custom programming.

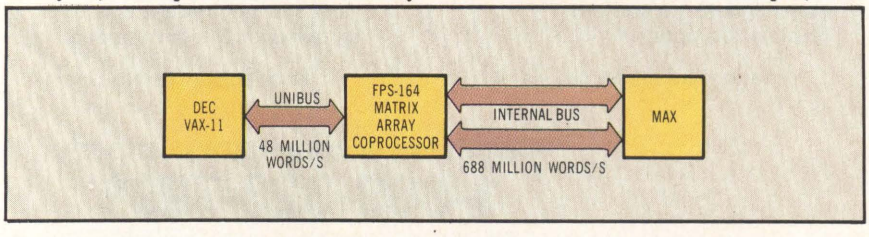
Design philosophy

The FPS-164/MAX, like its FPS-164 predecessor, is designed by Floating Point Systems (3601 SW Murray Blvd, Beaverton, OR 97005) to be a high speed Fortran engine. As such, it dedicates hardware logic to accelerate both floating point and matrix operations such as dot products, matrix multiplication, and matrix factoring. This makes them especially suitable for such areas as structural analysis, computational chemistry and physics, semiconductor physics, and molecular modeling.

However, this system does not act as a closely coupled coprocessor in the same way as an Intel 8087 floating point processor does with the 8086 microprocessor, according to Dr John Gustafson, one of the chief architects. Gustafson notes that the FPS-164/MAX does not act as a hardware "subroutine" summoned during execution of a program on the host processor. Rather, program development using the company's brand of Fortran-77 first occurs on the frontend processors (DEC VAX-11, IBM 308x/303x/4300, or Apollo Domain) with the compiled code then downloaded to the FPS-164/MAX for actual execution.

Furthermore, the attached processor has its own local memory of 56 million words (one word equals 64 bits) as well as auxiliary disk storage so that it does not rely on host resources. As a result, the matrix array coprocessor can run synchronously with programs executing on the host processor. This is because its performance is not limited to the bandwidth of the link between it and the host processor.

In this way, the attached processor looks more like a high speed peripheral device than a memory-mapped peripheral dependent on DMA transfers. For example, the bandwidth of the Unibus link between the FPS-164/MAX and a VAX-11 host is approximately 46 million words/s, while the attached processor's internal bus can transfer data at 688 million words/s. Since most memory transfers can be handled on its own bus, the system is relatively insensitive to the service interrupts that occur on Unibus from memory-mapped peripherals.



High throughput on the matrix array coprocessor (33 to 341 MFLOPS) is due to loose coupling between the processor and its frontend host, and tight coupling to the matrix accelerator portion. Relative throughputs are shown for a DEC VAX-11 configuration using a Unibus link.

Hardware architecture

Interrupt servicing becomes a performance bottleneck for other closely coupled coprocessor schemes because there is only one path between memory (both local and global) and the arithmetic processors, Gustafson says. This can be a problem for applications with a high degree of redundancy in either the operations performed or the data used. Examples include floating point operations such as fast Fourier transforms that have repetitious operations, or matrix operations like dot products that have redundant data.

The matrix array coprocessor seeks to overcome this bottleneck by creating as many local paths between arithmetic units and memory as possible. The MAX portion itself consists of from 4 to 30 identical arithmetic/memory cells that perform matrix operations in parallel. These cells are mapped into the top 1 million words of the 56-million word memory previously inaccessible to users. Since these arithmetic/memory cells look like "smart" memory to the control unit of the FPS-164/MAX, they can either be individually addressed (or addressed as a single unit) by Fortran or assembly language routines.

This scheme enhances parallel operations by allowing the bus to be used once to distribute a problem over several cells. In fact, vectors

(continued on page 40)



If you're a major league VAR, we want you on our team.

IBM is scouting for the most valuable of VARs: those with outstanding new ideas and a great batting average.

If you're one of them, you could become a Value Added Remarketer of IBM products. And what could that mean to you?

First, IBM can add clout to your marketing efforts. For example, we can help with product literature, direct mail and business show support. To add to your skills, IBM offers a wide range of professional classes for VARs.

Furthermore, thanks to the online referencing system used by our own sales force, we can direct prospects with special needs right to VARs with appropriate solutions.

And, as one of the finest of VARs, you'll be selling the finest equipment: some of IBM's most competitive products. Our VARs can apply for the IBM 4300 systems, System/36, Series/1, System/38, the IBM Personal Computer and the System 9000 family.

To find out more about the advantages of becoming an IBM VAR, simply send in the coupon below or call 1 800 IBM-VARS, Ext. 562.

If you think your company can qualify, now's the time to touch base.

IBM®

562/8-84

Larry Humphreys
IBM Distribution Channels
P.O. Box 76477
Atlanta, GA 30358

Please send me your free booklet, "Looking for Leaders."

Name_____ Title_____

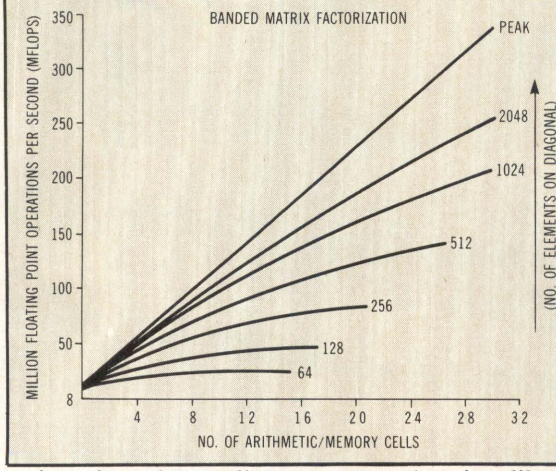
Company_____

Address_____

City_____ State_____ Zip_____

Phone_____ Product Interest_____

Matrix array processor
(continued from page 38)



The attached processor's performance depends on the number of arithmetic/memory cells configured, as well as the sparcity of the matrix being operated on (distribution of non-zero elements on the matrix diagonal). Higher performance is possible when there are many elements on the diagonal itself.

written into these cells are processed synchronously and in parallel with other operations occurring in the remainder of the system's memory. As a result, as many as 124 vector dot products can be computed simultaneously.

In addition, this memory-mapped technique keeps setup and postprocessing overhead low. Thus, even operations involving short vectors can take advantage of MAX with minimum penalty. Furthermore, Gustafson notes that the matrix array coprocessor reduces the time needed to process a system of equations so that it equals the time spent in setup and postprocessing. Present implementations have about 30 percent of the execution time spent in setup and postprocessing, with the remaining 70 percent involving the actual processing of the system of equations.

As a result, the system is fast even when compared to a VAX-11/FPS-164 operating in tandem. Gustafson has run a benchmark that utilizes an earthquake analysis of a building in San Francisco. The VAX-11/FPS-164 combination performed the analysis at the rate of 7.5 million floating point operations per second (MFLOPS). In contrast, a similar configuration with MAX capability ran the same code at 20 MFLOPS. A VAX-11 executing the code without the benefit of either ran the code 46 times slower than the VAX-11/FPS-164 benchmark.

Programmers need not worry about these varying performance levels if they use the fast matrix solution library (FMSLIB) provided with the system. The library contains a package of matrix algebra routines written in FPS-164 assembly language that can be incorporated into new or existing Fortran programs. The routines can handle either real or complex data with the ability to efficiently manipulate a variety of sparcity patterns (the distribution of non-zero elements along the diagonal of the matrix).

Programmers wishing to access the system's resources can use a Fortran-77 cross compiler that runs on the host machine, or an assembler that generates binary object modules. The generated code for the Fortran cross compiler is optimized for use on the attached processor's hardware architecture, and includes extensions to access the FPS-164 disk subsystem, host files, and the user's host terminal.

Since the system can run in a standalone fashion, a system job executive is also available to manage permanent and scratch files, support Fortran I/O operations, and track runtime overlays. The executive can accept commands from either interactive foreground sessions on the processor itself or from background batch job streams on the host computer. Users can attach/detach the processor from the host and copy files between the two in addition to executing complete jobs on the system itself.

—Joseph Aseo, Field Editor
SYSTEM TECHNOLOGY
(continued on page 44)

Potter & Brumfield Authorized Sales Representatives:

Central Region

IL, Elk Grove Village
Oasis Sales Corporation
Phone: 312/640-1850

IN, Fort Wayne
R.O. Whitesell & Assoc., Inc.
Phone: 219/432-5591

IN, Indianapolis
R.O. Whitesell & Assoc., Inc.
Phone: 317/359-9283

IN, Kokomo
R. O. Whitesell & Assoc., Inc.
Phone: 317/457-9127

KY, Louisville
R. O. Whitesell & Assoc., Inc.
Phone: 502/426-7696

MI, Grand Rapids
R. O. Whitesell & Assoc., Inc.
Phone: 616/942-5420

MI, Southfield
R. O. Whitesell & Assoc., Inc.
Phone: 313/559-5454

MI, St. Joseph
R. O. Whitesell & Assoc., Inc.
Phone: 616/983-7337

MN, Edina
Mel Foster Technical Sales, Inc.
Phone: 612/941-9790

OH, Cincinnati
R. O. Whitesell & Assoc., Inc.
Phone: 513/521-2290

OH, Cleveland
R. O. Whitesell & Assoc., Inc.
Phone: 216/447-9020

OH, Columbus
R. O. Whitesell & Assoc., Inc.
Phone: 614/888-9396

OH, Dayton
R. O. Whitesell & Assoc., Inc.
Phone: 513/298-9546

PA, Pittsburgh
R. O. Whitesell & Assoc., Inc.
Phone: 412/963-6161

WI, Milwaukee
E. A. Dickinson & Assoc., Inc.
Phone: 414/264-1080

Eastern Region

MD, Baltimore
Micro-Comp, Inc.
Phone: 301/247-0400

MA, Lexington
Hathaway Electronics, Inc.
Phone: 617/861-7010

NJ, Medford
Jack W. McCoy, Inc.
Phone: 609/953-0770

NY, Melville
Comtronic Associates, Inc.
Phone: 516/249-0505

NY, Rochester
Ossmann Component
Sales Corp.
Phone: 716/424-4460

NY, Syracuse
Ossmann Component
Sales Corp.
Phone: 315/455-6611

PA, Harrisburg
Beil & Whitaker, Inc.
Phone: 717/564-6900

PA, Orefield
Beil & Whitaker, Inc.
Phone: 215/395-0150

PA, Reading
Beil & Whitaker, Inc.
Phone: 215/779-2610

PA, Wernersville
Beil & Whitaker, Inc.
Phone: 215/670-1111

Puerto Rico, San Juan
Electronic Technical Sales, Inc.
Phone: 809/790-1300

Western Region

AZ, Phoenix
Howe & Howe Sales, Inc.
Phone: 602/264-7971

CA, Culver City
Bestronics, Inc.
Phone: 213/870-9191

CA, Irvine
Bestronics, Inc.
Phone: 714/261-7233

CA, Los Altos
Elliott Recht Associates
Phone: 415/964-6321

CA, San Diego
Bestronics, Inc.
Phone: 714/452-5500

CA, Woodland Hills
Bestronics, Inc.
Phone: 818/704-5616

CO, Littleton
Simpson Associates, Inc.
Phone: 303/794-8381

NM, Albuquerque
C.T. Carlberg & Assoc., Inc.
Phone: 505/888-3883

UT, Midvale
Simpson Associates, Inc.
Phone: 801/566-3691

WA, Seattle
Northmar, Inc.
Phone: 206/524-5170

Southern Region

AL, Huntsville
Cartwright & Bean, Inc.
Phone: 205/830-1540

FL, Orlando
Cartwright & Bean, Inc.
Phone: 305/422-4531

FL, Fort Lauderdale
Cartwright & Bean, Inc.
Phone: 305/735-4900

GA, Atlanta
Cartwright & Bean, Inc.
Phone: 404/233-2939

IA, Cedar Rapids
P-M-R, Inc.
Phone: 319/362-9177

KS, Overland Park
P-M-R, Inc.
Phone: 913/381-0004

KS, Wichita
P-M-R, Inc.
Phone: 316/684-4141

LA, Kenner
Cartwright & Bean, Inc.
Phone: 504/466-0020

MS, Jackson
Cartwright & Bean, Inc.
Phone: 601/981-1170

MO, Maryland Heights
P-M-R, Inc.
Phone: 314/569-1220

NC, Charlotte
Cartwright & Bean, Inc.
Phone: 704/377-5673

NC, Raleigh
Cartwright & Bean, Inc.
Phone: 919/781-6560

OK, Tulsa
Components, Inc.
Phone: 918/583-9149

TN, Knoxville
Cartwright & Bean, Inc.
Phone: 615/693-7450

TN, Memphis
Cartwright & Bean, Inc.
Phone: 901/276-4442

TX, Austin
Ammon & Rizos Company
Phone: 512/454-5131

TX, Dallas
Ammon & Rizos Company
Phone: 214/233-5591

TX, Houston
Ammon & Rizos Company
Phone: 713/781-6240

AMF
Potter & Brumfield



P&B... the power house of the 80's.

For switching power supplies with the technology and reliability your design requires, rely on Potter & Brumfield. With U.S.-based design and manufacturing functions, we can provide the affordable, high quality switchers you need, when you need them.

Subsystems Experience

We're primarily known as a component manufacturer, but for nearly 20 years we have been designing and building electronic subassemblies on a custom basis for major OEMs. This expertise in electronics is now being committed to a standard line of switching power supplies.

Meeting Your Needs

Experienced design engineers and assembly technicians, with the aid of CAD/CAM, optimize the performance of our standard switchers for your particular application. Automated manufacturing equipment helps assure consistently high quality, and computer-based testing of each unit helps us meet our goal of 100% customer satisfaction.

Standard Models Available

Our standard open-frame models incorporate an efficient, MOSFET design. As many as five outputs are offered on 100 - 300 watt units. All are certified to meet FCC and VDE requirements for conducted and radiated EMI/RFI, as well as safety standards established by UL, CSA, IEC and VDE.



Discover Our Capabilities

Find out more about switching power supplies from Potter & Brumfield. Call one of our sales offices and review your needs, or call toll-free 800/223-1416 for the name of your nearest sales representative or authorized distributor. Potter & Brumfield Division AMF Incorporated, 200 Richland Creek Drive, Princeton, IN 47671.

Regional Sales: Oak Brook, IL, 312/887-0811; Braintree, MA, 617/848-6550; San Juan Capistrano, CA, 714/493-4503; Norcross, GA, 404/449-4601.

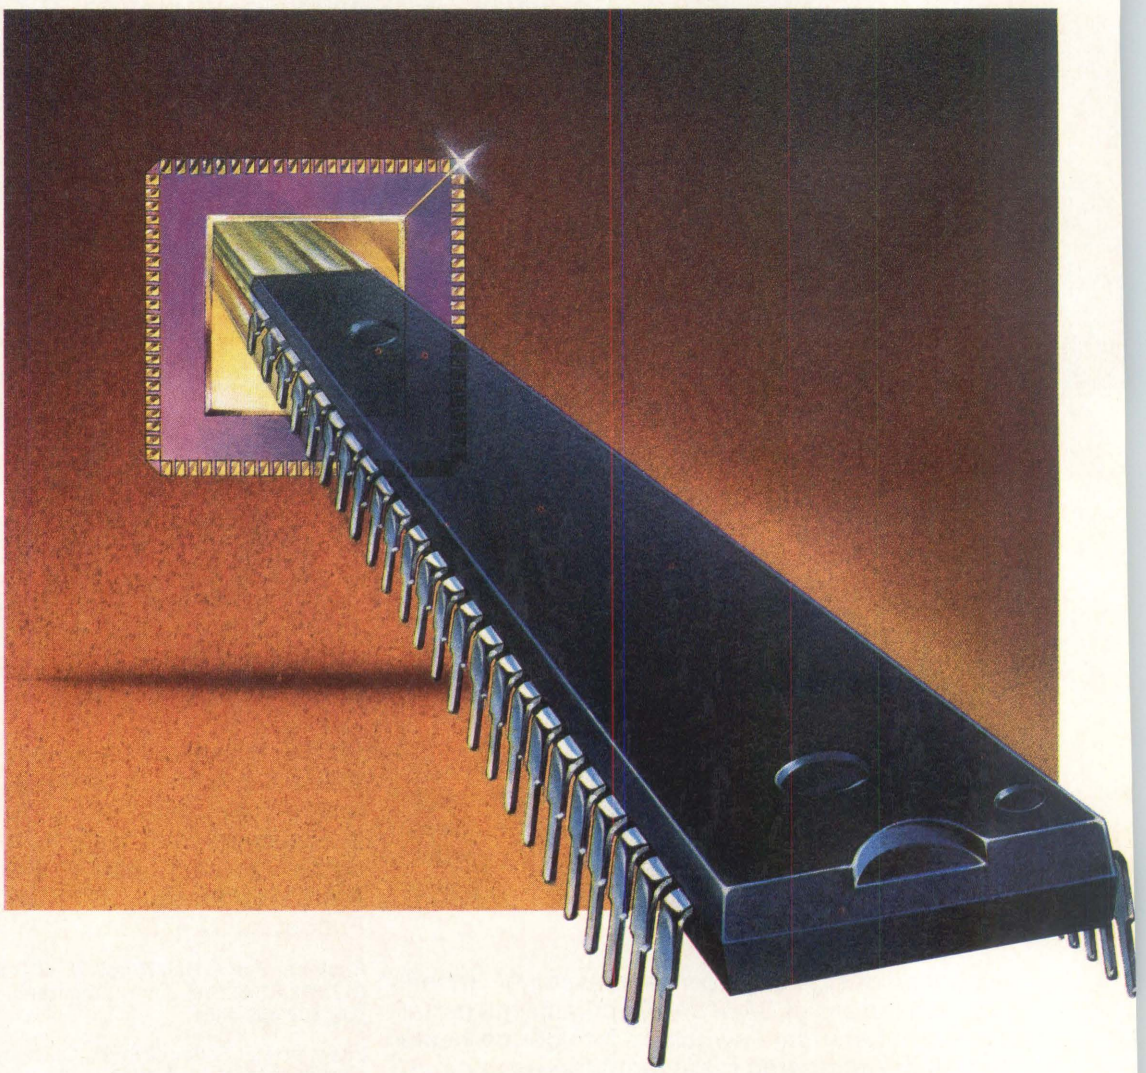
International Sales: Guelph, Ontario, 519/822-1576, TELEX: 0695-6522; Bristol, England, (0272) 716301; Yokohama, Japan, (045) 812-1418; Hong Kong, TELEX: 39555; Singapore, TELEX: RS 34686; Chatswood, N.S.W. Australia, (02) 411-5222.

AMF

Potter & Brumfield

CIRCLE 23

DOUBL



E DUTY

Introducing the MK68200 16-bit microcomputer/microprocessor.

Unlike any other 8- or 16-bit single-chip microcomputer, the new Mostek MK68200 performs equally well as a single-chip microcontroller and as a peripheral I/O controller. In the latter type of application, the MK68200 can provide local intelligence required to handle typical computer system I/O functions. Without tying up the host processor. In addition, it can transfer data to and from system memory using a software DMA function. So, as an example, it's ideal as a front-end processor for handling complex serial I/O protocols.

Actually, the 16-bit MK68200 gives you a whole new dimension of system control. One that enables you to monitor, measure, regulate, sense, and interconnect with more precision than ever before.

The operative word is more. More speed. More powerful instructions. More I/O functions. More efficiency. More versatility. All of which mean more application opportunities. For robotics. Engine control. Pattern recognition. And real-time measurements and control – as a stand-alone device, or in tandem with other microprocessors as an intelligent peripheral controller.

First, consider speed. Using the 6 MHz instruction clock rate, the MK68200 can execute a 16-bit multiply as fast as a 3.5μs; a 32 x 16 divide in 3.8μs; and a 16-bit add/subtract in as little as 500ns. In other words, it's faster than any other general-purpose, single-chip micro currently available.

One reason for this faster speed is code efficiency; most MK68200 instructions are just one word. That not only saves code space, it also improves execution speed. And the powerful instruction set incorporates more than 50 instruction types which operate on both byte and word operands.

Next, look at the available versions. The MK68200 comes in a 48-pin DIP with 0 or 4K bytes of ROM and 256 bytes of RAM. It is expandable to address a full 64K byte address space externally. And finally, an 84-pin LCC version (with no ROM), is available for system development or for multiple bus applications.

For single-chip applications, up to 40 pins of the 48 pins are available for I/O to include two 16-bit parallel ports. Plus a full-duplex USART with double-buffered transmit and receive capable of operating at data rates up to 1.5Mbps (using a 6MHz instruction clock).

The MK68200 offers extensive interrupt capabilities. One non-maskable interrupt input is provided as well as 14 maskable interrupt sources. In addition, all 14 independently vectored interrupts are user-assignable in software for any priority scheme. And to enhance serial communications, there's also an innovative wake-up interrupt on the serial channel.

As for support, there's a powerful Macro Cross Assembler, along with a host of other software tools. As well as RADIUS™, a very cost-efficient remote development station. And to expand the MK68200 family, plans include higher performance versions with denser memory and faster throughput. Plus power-conserving CMOS.

Look beyond the limits of conventional single-chip systems into the realm of ultimate control. It's yours with the MK68200. For more information, contact Mostek, 1215 W. Crosby Road, MS2205, Carrollton, TX 75006, (214) 466-6000. In Europe, (32) 02/762.18.80. In Japan, 03/496-4221. In the Far East (Hong Kong), 5-681157-9.

RADIUS is a trademark of Mostek Corporation.



CIRCLE 24

Big Blue makes a major move to link computers

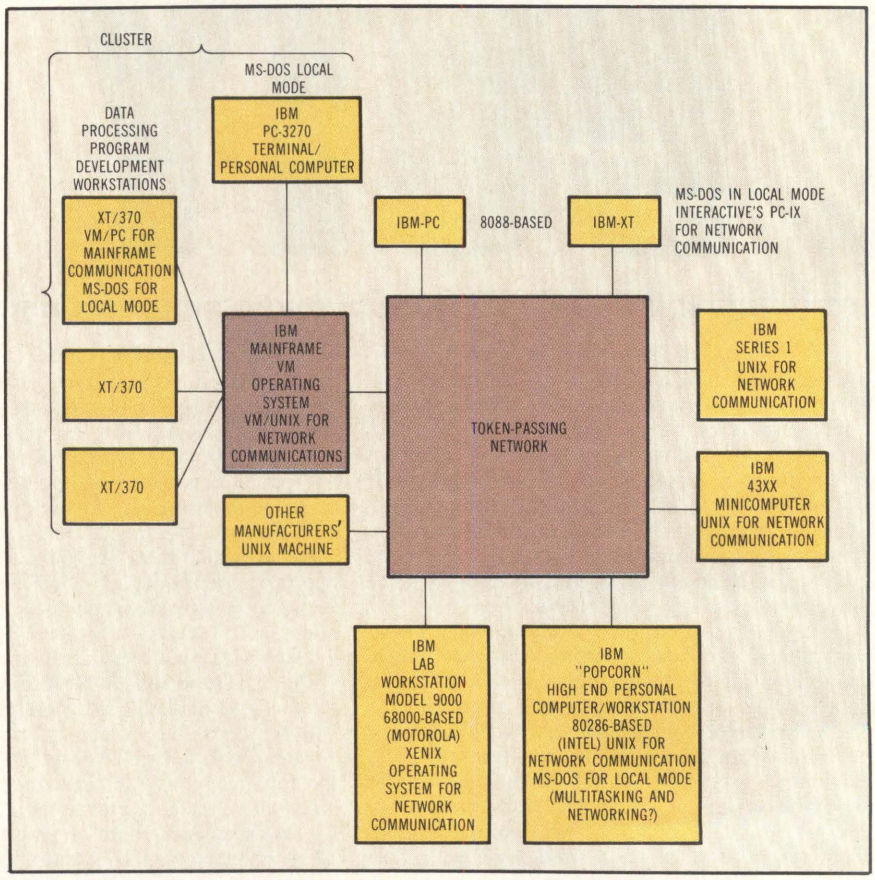
Code-named VM/IX, this IBM licensed version of AT&T's Unix operating system clarifies the line of march that IBM is taking toward networking workstations, terminals, microcomputers, minicomputers, and mainframes. Operating as tasks under IBM's interactive VM mainframe operating system, the VM/IX operating systems serve as IBM's next step on the way to accomplishing its networking goals. This was the sentiment of Unix community insiders attending USENIX, the Unix operating system user's group meeting held in Salt Lake City, Utah, last June.

The VM/IX operating system and specific machine-tailored versions will eventually operate on IBM mainframes that run VM, as well as minicomputers such as the 43XX series, Series/1 machines, the XT/370, and the IBM XT itself. In short, computer system designers will be able to link such machines in Unix-controlled communication networks while the machine's native operating system does specific computing chores.

VM/IX is the result of IBM's work with Interactive Systems Corp. Interactive developed the single-user Unix for the XT that became IBM's PC-IX operating system. VM/IX is right in line with IBM's long-term goals. Currently, IBM dominates the mainframe market. Thus, the first goal is to make sure that other computers/workstations that are directly connected to an IBM mainframe remain IBM machines running IBM operating systems. The VM/IX will help ensure that this goal is met.

The second goal is to ensure that IBM minicomputers grow in market share [Digital Equipment Corp (Maynard, Mass) dominates the minicomputer market], and that IBM micros continue to be best-sellers. IBM feels that improving the product, cutting the price, and allowing networking through the use of VM/IX builds the strategy for achieving this goal.

The Unix operating system was chosen for this application because it can run on almost any machine. It is portable because of its design and its C-code, which has minimal hardware dependencies. (C is a language midway between assembler and a high level language.)



The communication link hooking all the departments in the corporate environment will feature a Unix-based communication system that links IBM and other vendors' machines, regardless of type. Both local and remote machine operation will be possible. Each machine will be able to access data or programs from the others and will perform either general purpose or dedicated functions.

In addition to its utilities and program development tools, Unix is the only portable operating system with so much experience and exposure. Thus, while most operating systems are proprietary, Unix can run on any manufacturer's minicomputer. Moreover, it allows system designers to link, not only IBM minis, but products from other companies such as DEC, Data General, Hewlett-Packard, and Gould.

As a bonus, the system designer can even link IBM's Motorola 68000-based 9000 Series scientific minicomputer/workstation. Because it already runs on Microsoft Corp's (Bellevue, Wash) Unix-licensed Xenix operating system, conversion would be straightforward.

Versatility and development

This common VM-based operating system can serve as the basis for network communication either using a peer-to-peer or master-slave variation of IBM's systems network architecture (SNA) architecture or the International Standards Organization (ISO) seven-layer model for computer communications. In some cases, the software that takes care of these protocol functions will run on the operating system in the computer. In other cases, for the sake of low overhead, it will run on a frontend processor. And, for those just needing a personal computer/workstation and a

(continued on page 46)

68000
16000
8086/88
Z8000

SOFTWARE DEVELOPMENT TOOLS

"One-Stop Shopping"

OASYS provides a **"One-Stop Shopping"** service for software developers and managers in need of proven, cost effective, **cross-** and **native-** development tools.

OASYS can save you time, energy and money! We understand what it means to be a developer. Over the past 3 years, we've built over 1MB of working code.

We not only develop our own tools, but also specialize in evaluating, selecting and distributing the best complementary tools from other suppliers.

Our tools are currently in use in over **1,000** installations worldwide on **micro-, mini-,** and **mainframe** computers for a variety of **8-, 16-** and **32-** bit UNIX (and non-UNIX) systems.

Most likely, we have what you're looking for (even if it doesn't appear in the tables shown). But, if we don't, we'll be glad to tell you who does.

So, call or write today for more information and start shopping the smart way, the fast way, the economical way.

"The One-Stop Shopping Way."

OASYS

60 ABERDEEN AVENUE
CAMBRIDGE, MA 02138
(617) 491-4180

CROSS TOOLS

PRODUCTS (1)	HOST (2)	TARGET (3)
C COMPILERS	VAX, PRIME	68000 16000 8086/88
PASCAL COMPILERS	VAX PDP-11, LSI-11 PRIME	68000 16000 8086/88
FORTRAN COMPILERS	VAX PDP-11, LSI-11	68000 16000 8086/88
ASSEMBLERS (4)	VAX, PDP-11, LSI-11, PRIME, IBM/PC, IBM 370	68000, 16000, 8086/88, Z8000, 680X, 808X, Z80
SIMULATORS	VAX, PDP-11 LSI-11, PRIME, IBM/PC, IBM 370	68000, 8086/88 808X, Z80

(1) WE DISTRIBUTE PRODUCTS FOR: GREEN HILLS SOFTWARE, VIRTUAL SYSTEMS, COMPLETE SOFTWARE, PACER SOFTWARE; SOFTWARE MANUFACTURERS
(2) HOST OPERATING SYSTEMS INCLUDE: VMS, RSX, RT-11, PRIMOS, UNIX V7, III, V, BSD 4.1, 4.2, UNOS, IDRIS, XENIX, MS/DOS, VM/CMS, CPM 68K
(3) OTHER TARGETS ARE: M6801-6803, 6805, 6809, 8080, 85, 28, 35, 48, 51; Z-80
(4) ALL ASSEMBLERS INCLUDE LINKER, LIBRARIAN AND CROSS-REFERENCE FACILITY
(5) AVAILABLE ON: CALLAN, OMNIBYTE, CHARLES RIVER DATA, PLEXUS, SAGE, FORTUNE, WICAT . . . to name a few.

C/UNIX NATIVE TOOLS

- **NATIVE ASSEMBLERS FOR 68000s** (4, 5)
- **SYMBOLIC C SOURCE CODE DEBUGGER**
- **C-TIME PERFORMANCE UTILITY**
- **UP/DOWN LINE LOAD UTILITIES**
- **COMMUNICATION UTILITIES**
- **C-BASED FLOATING POINT MATH PACKAGE**
- **BASIC-TO-C TRANSLATOR** . . . AND MORE

TRADEMARKS: UNIX IS A TRADEMARK OF BELL LABORATORIES, XENIX AND MS/DOS ARE MICROSOFT CORP'S; IBM/PC, VM/CMS, AND IBM 370 ARE INT'L BUSINESS MACHINES; VAX, PDP-11, LSI-11, VMS, RSX, AND RT-11 ARE TRADEMARKS OF DIGITAL EQUIPMENT CORP; CPM 68K IS DIGITAL RESEARCH'S; PRIMOS IS PRIME'S; UNOS IS CHARLES RIVER DATA'S; IDRIS IS WHITESMITH'S LTD.

Big Blue's move

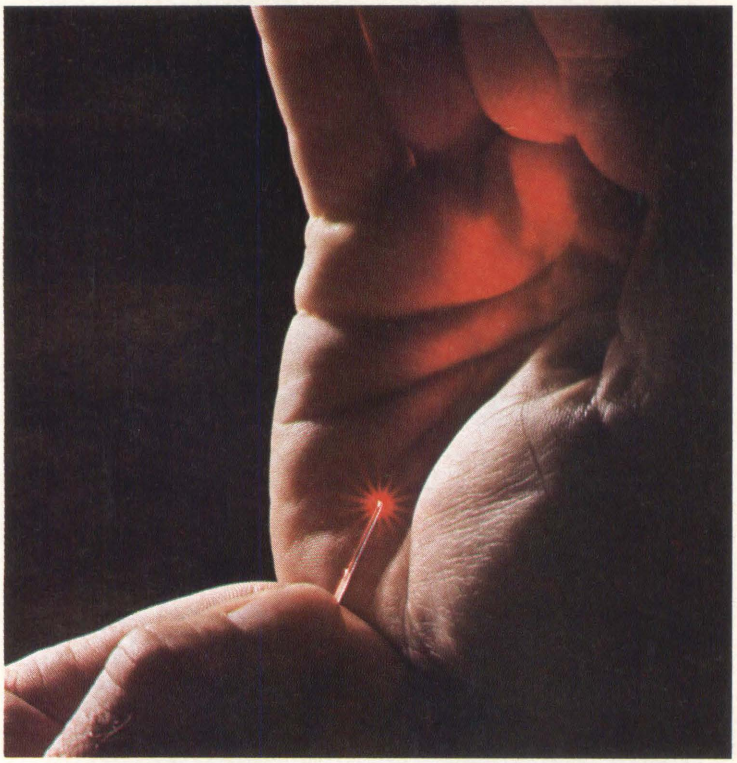
(continued from page 44)

mainframe link for program development, there is IBM's VM/CMS approach, which will be network-linkable through the mainframe.

VM/CMS links the program development/personal computer XT/370 to a mainframe through a common operating system. This is done without using Unix, whose capabilities are not needed for this application. The XT/370 was introduced mainly for local development of mainframe programs under VM. But, it also allows the use of MS-DOS-based personal productivity software in a local mode.

BREAKTHROUGH!

Totally new optical fibers like these handle data communications in computer installations far better than wire. (And the cables are as easy to handle as wire coax.)



Ensign-Bickford's new proprietary hard cladding chemically bonds to the silica core. Resultant fibers are so tough they handle like wire, install quickly and reliably, provide long, EMI-resistant life. Large core and high NA mean easy coupling. High strength and microbend resistance mean easy deployment—and you can use standard, off-the-shelf connectors!

Call or write for the data you need to specify from a wide range of fiber, cable and terminated assemblies.

Ensign-Bickford Optics company

a subsidiary of Ensign-Bickford industries, inc.

660 Hopmeadow St. Simsbury, Conn. 06070 203/658-4411

IBM moved portions of its VM/CMS operating system to the XT/370 so it could become the micro part of a micro-to-mainframe link's common operating system. This link, specifically designed for professional program developers of mainframe and minicomputer software, is geared to the mainframe/XT/370 subnetwork (cluster).

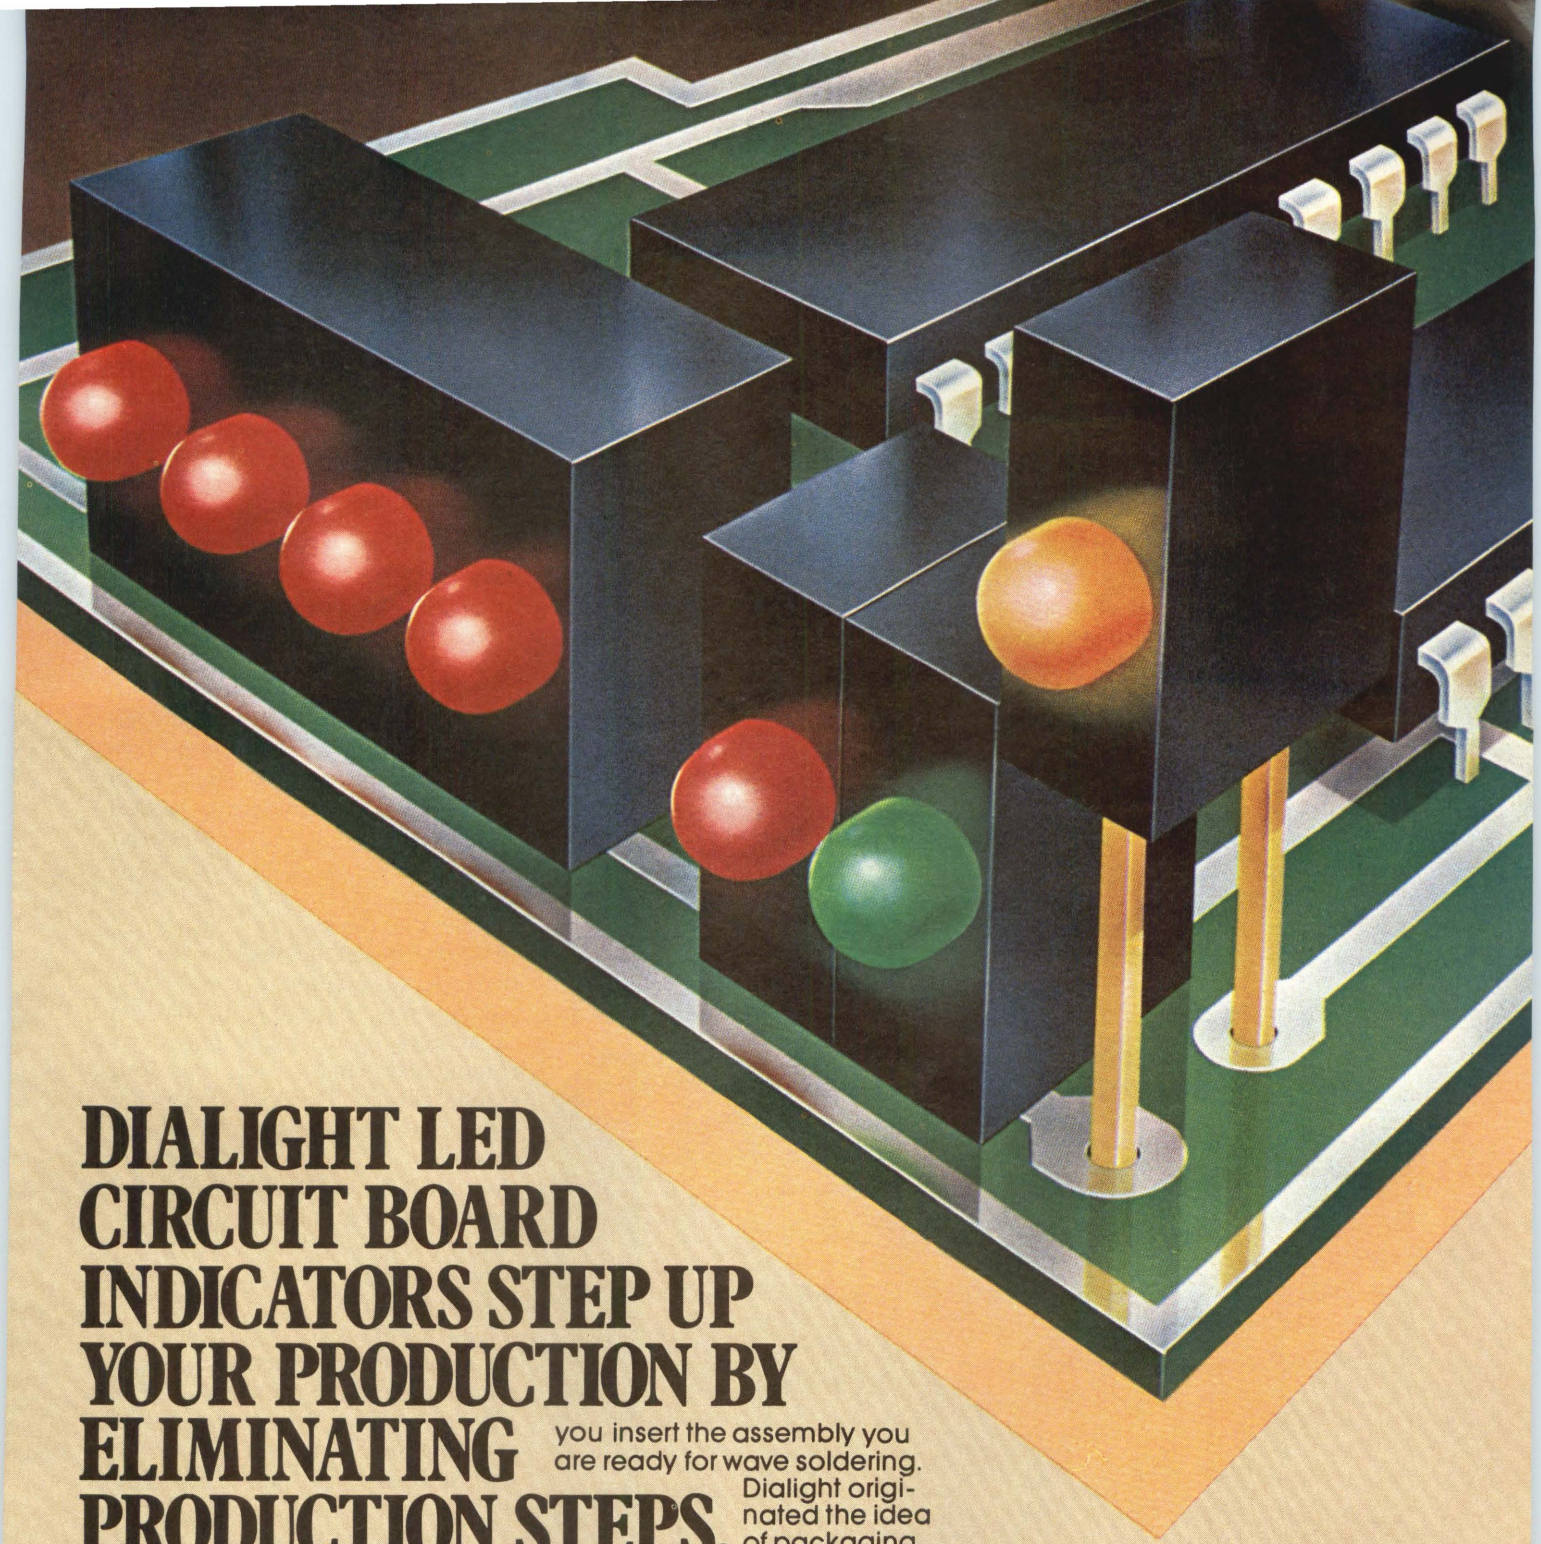
In part, IBM introduced the VM/CMS operating system because it already exists on mainframes. The existence of the same operating system on both sides of a micro-to-mainframe link facilitates file access, information transfer, and program execution across the connection. A cluster of XT/370s operates as a standalone or communicates with networked computers through a mainframe.

VM/CMS offers advantages on the mainframe side of the link also. Here, it is used for data processing, software program development, and, many management information systems departments in major firms already know it. This saves learning time as well.

VM/CMS will run in standalone mode on the 4301 and 4311 machines—tying IBM's neat computer network world together even more tightly. These high end, desktop superminicomputers allow all the functions of such machines for the computer designer. Designers must also realize that upcoming engineering and scientific applications (eg, on the 43XX series) will cut into workstation applications heretofore dominated by firms like DEC and Data General.

The soon-to-be-announced, multitasking, high end personal computer/workstation, code named Popcorn, will also be part of IBM's contribution to these applications. This new machine will use Unix for networking (its Intel 80286 has memory management capabilities), and will have a version of MS-DOS that will handle either multitasking, networking, or both for local mode, personal productivity operations. If IBM has Popcorn run VM/CMS-based software, and has the printed circuit boards

(continued on page 48)



DIALIGHT LED CIRCUIT BOARD INDICATORS STEP UP YOUR PRODUCTION BY ELIMINATING PRODUCTION STEPS.

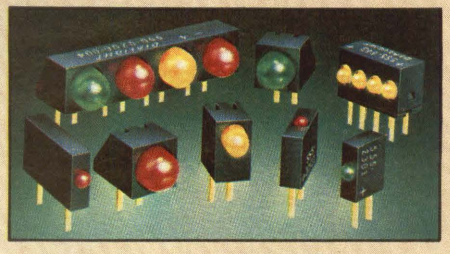
You'll save money when you stop mounting LEDs on PC boards the old way — bending leads, inserting holders, adding resistors — and start using LED Circuit Board Indicators from Dialight.

Mounting our LED Circuit Board Indicators is easier and less time-consuming. They eliminate production steps and reduce labor costs. Not only is positioning faster, it's far more accurate. As soon as you insert the assembly you are ready for wave soldering.

Dialight originated the idea of packaging LEDs for easy mounting on PC boards. And we've developed over 50 different Circuit Board Indicators in red, green, yellow and red/green bicolor. Choose single-element LEDs or QUAD-LED™ four-element arrays with a wide range of voltages with or without current limiting resistors.

Send for our catalog. And the next time you need LEDs for PC boards, eliminate steps and save money — specify Dialight. 203 Harrison Place, Bklyn., NY 11237 (212) 497-7600 TWX: 710-584-5487

See us at Midcon, Booth #930



DIALIGHT
meets your needs.
A North American Philips Company

CIRCLE 27

Big Blue's move

(continued from page 46)

(which define the XT/370) incorporated as Popcorn add-ons, computer system designers will have more flexibility. They will also be able to provide customers with a variety of minicomputer and mainframe features on a personal machine.

Behind the design

As designed, the XT/370 has two operating systems—the most recent version of PC-DOS, and IBM's VM/PC, which is a subset of mainframe-based CMS. The theory behind this design is that VM supplies mainframe users with virtual terminals. It works with an operating system like CMS to execute an application program. But, a dedicated microcomputer like the XT/370 does not need software to set up virtual terminals or the CMS parts that take care of punched card inputs to mainframes.

Even though these software functions were not incorporated into the XT/370, the remaining software would have been crowded on a 10-Mbyte hard disk. Thus, the XT/370 is equipped to swap software between its RAM, hard disk, and mainframe counterpart.

Such swapping is usually slow and inefficient. To avoid these problems, the XT/370 has three microprocessors that operate all VM/CMS mainframe software, as well as the entire IBM 370 instruction set. With one IBM proprietary 68000 processor and one standard 68000 processor, plus an Intel 8087 floating point processor, the machine is far more advanced than the IBM PC or XT, which have only an Intel 8088 processor.

In addition to the processors, this machine has a special control program that makes up for CMS needing a VM environment, and allows MS-DOS to do I/O chores in collaboration with CMS. Of course, MS-DOS is still able to run the application programs designed for it when the XT/370 is in its local mode.

Thus, the XT/370 will be a link with IBM mainframe operating systems and software, and a local, personal computer that runs application programs such as Lotus 1-2-3, WordStar, and VisiCalc. To encourage the success of this venture, IBM is prompting third-party vendors to supply PC software that runs under VM/CMS. Expect IBM to encourage VM/CMS software for its minis, and maybe even for the Popcorn.

But, IBM is not the only company using VM/CMS. Several other manufacturers of micro-to-mainframe links (there are about a dozen announced products) are porting their software to work under this operating system. This means that once data is downline-loaded from the mainframe to the XT/370, it can be directly transferred from the VM operating system to the PC-DOS operating system application. The XT/370, by design, can look for and read mainframe directories and access files, downloading them to the PC-DOS operating system.

Ultimately, the success or failure of these networking plans hinges on computer system designers' operating system choice. This choice depends largely on IBM's ability to persuade designers to opt for its Unix rather than AT&T's. AT&T will counter this pressure with arguments of its own. In any event, it remains a guaranteed winner whatever the choice because it licenses Unix.

The most persuasive argument may turn out to be experience. Many designers prefer to obtain all of their software from one source, thus avoiding the blame shifting typically encountered when something goes wrong with multivendor systems.

Understandably, IBM will be reluctant to port its Unix to competitive machines. Thus, AT&T or any of the Unix software houses that have recently emerged will likely supply Unix for other vendors' equipment within the IBM network.

—Harvey J. Hindin,
Special Features Editor

SYSTEM TECHNOLOGY

(continued on page 53)

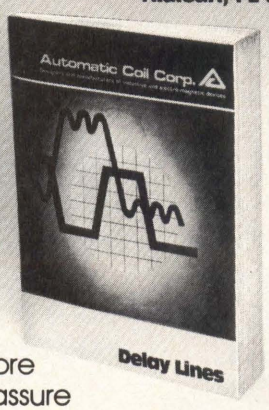
NEW JUST RELEASED! FREE

Delay Line Catalog Features Detailed Technical Data Section

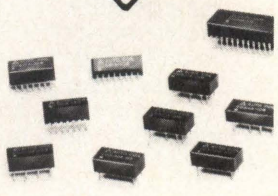
Operating characteristics, design and usage are thoroughly discussed in this 32 page catalog (AC-37). Formulas and graphs cover delay time, impedance, rise time, bandwidth, figure of merit, attenuation and phase compensation.

The product section presents over 650 standardized active and passive delay lines including TTL digital and function modules. Manufacturing is at our off-shore Carribean plants to assure you lowest pricing. Request your copy today by calling (305) 883-5311 or contacting:

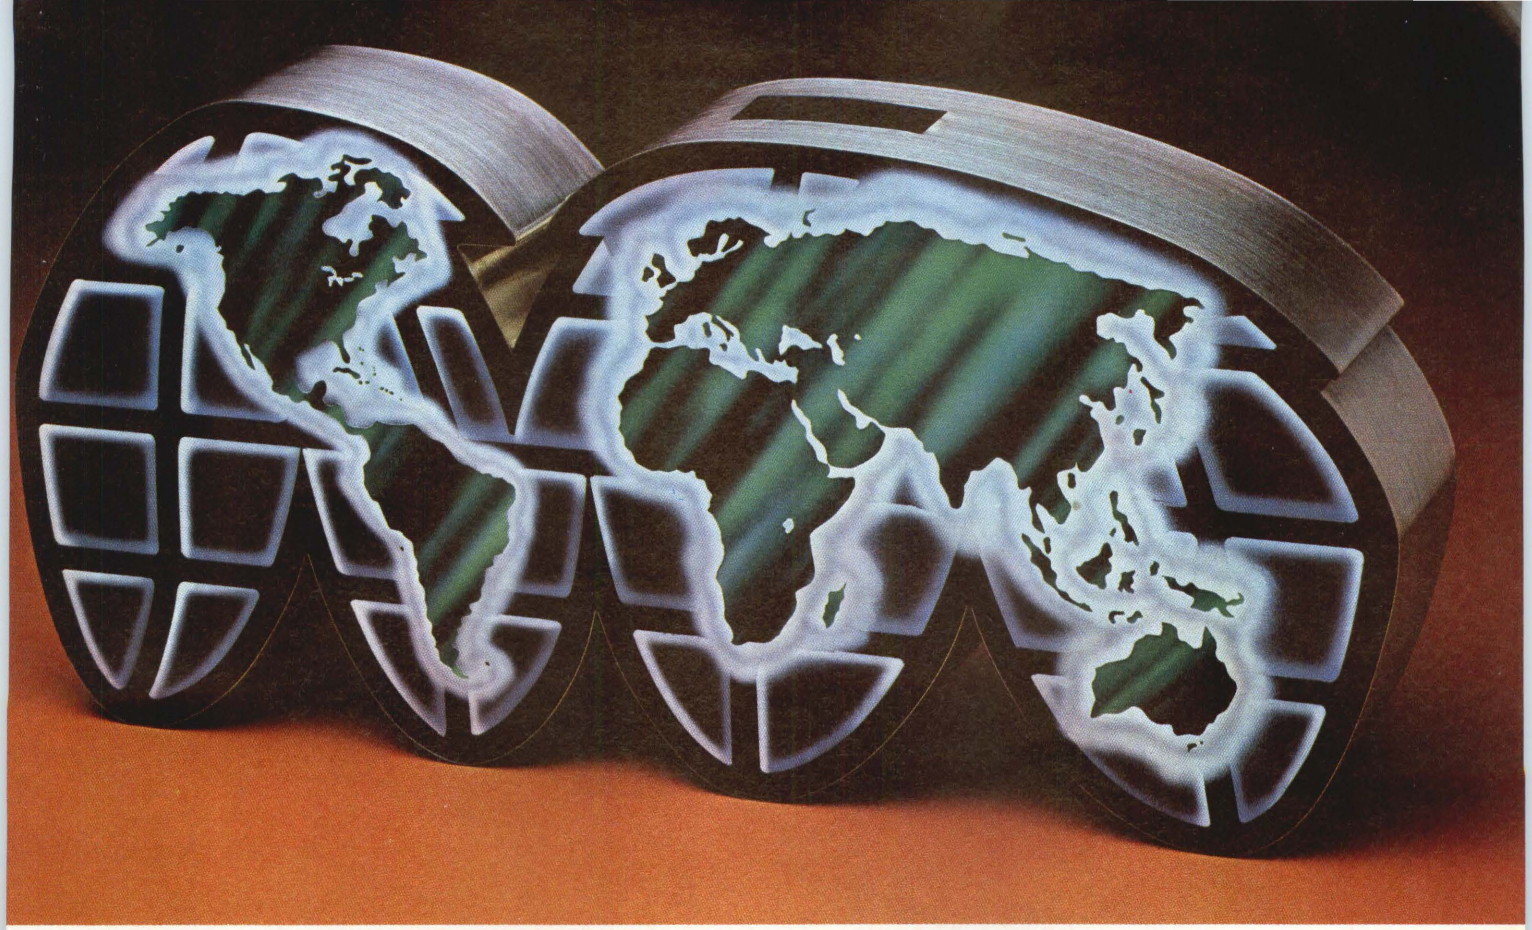
Automatic Coil Corp.
Subsidiary of Designatronics, Inc.
455 East 10th Avenue
Hialeah, FL 33010



Smallest Active SIP Made!



CIRCLE 28



The Strategic Supplier Is One Who Shares The Risks As Well As The Rewards.

Capital-intensive program development is the price of the innovative product that takes the market by storm. The more external resources available to your company, the shorter development lead time, and the quicker return on investment.

Even in an erratic economy, GE Plastics continued to invest in facilities and people that are ready to help expand your own technology and opportunities. The $25 million Plastics Technology Center now nearing completion at the Plastics Group's global headquarters is just one example of major new application development and manufacturing sites worldwide, and of the demonstrated commitment to an active partnership role in enhancing your innovations.

The strategic supplier is one who commits tangible support to gain mutual growth. The strategic supplier is General Electric Plastics.

LEXAN® sheet/film/foam/resin
NORYL® foam/resin
VALOX® film/foam/resin
ULTEM® foam/resin
XENOY® resin

We bring good things to life.

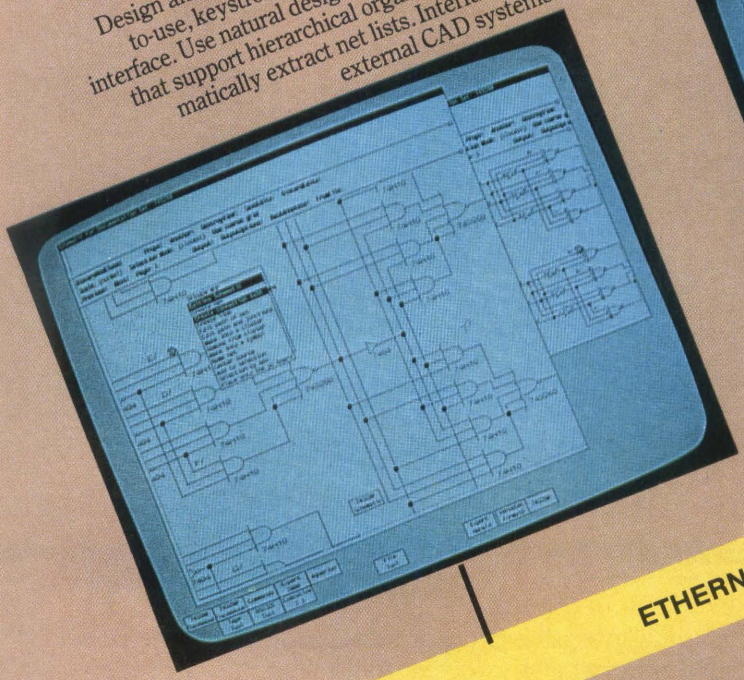
GENERAL ELECTRIC

® Registered Trademarks of General Electric Company

CIRCLE 29

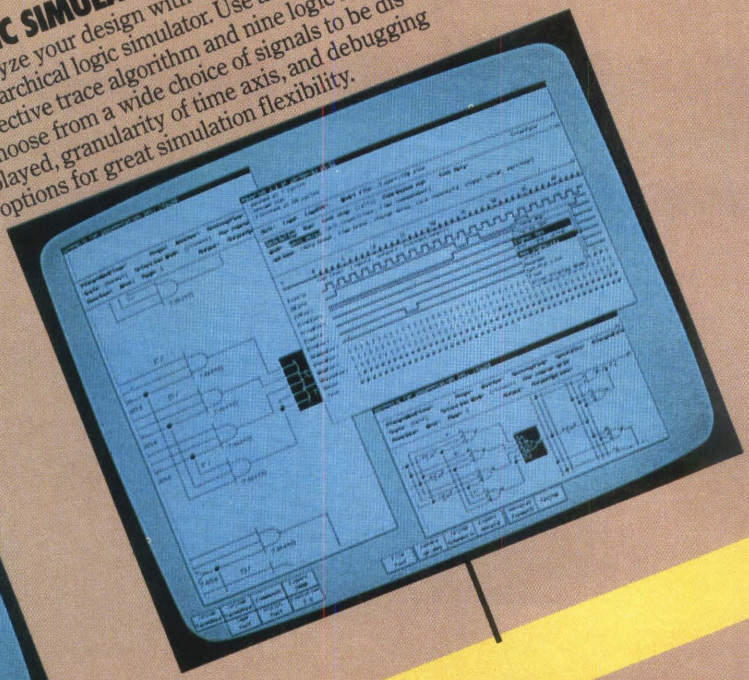
SCHEMATIC DESIGN

Design and edit logic diagrams with an easy-to-use, keystroke-efficient, interactive user interface. Use natural design capture techniques that support hierarchical organization. Automatically extract net lists. Interface with external CAD systems.

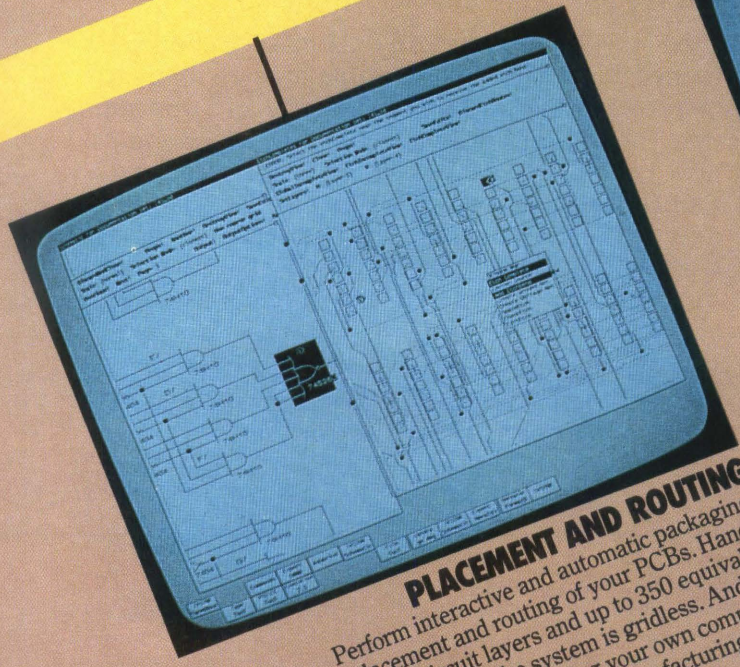


LOGIC SIMULATION

Analyze your design with an interactive hierarchical logic simulator. Use an event-driven selective trace algorithm and nine logic states. Choose from a wide choice of signals to be displayed, granularity of time axis, and debugging options for great simulation flexibility.

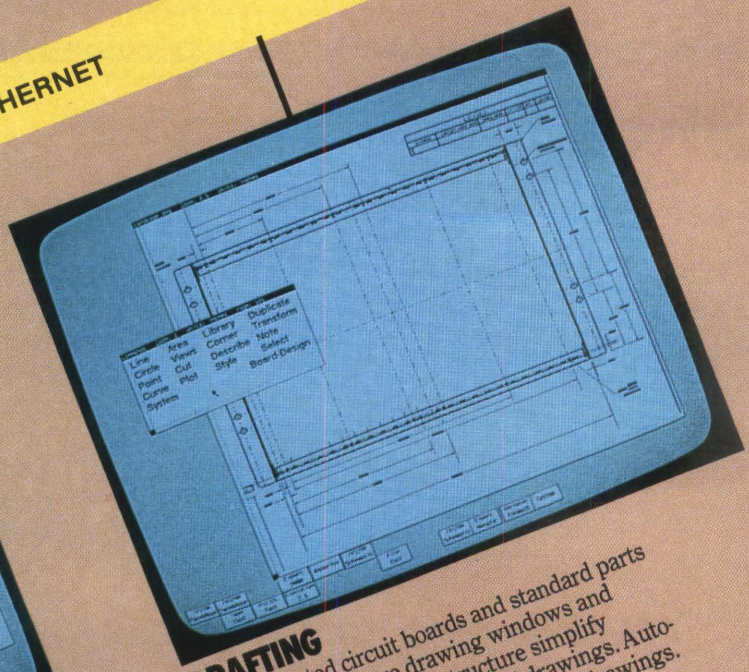


ETHERNET



PLACEMENT AND ROUTING

Perform interactive and automatic packaging, placement and routing of your PCBs. Handle up to 16 circuit layers and up to 350 equivalent ICs. Best of all, the system is gridless. And you can design boards using your own company's manufacturing rules.



DRAFTING

Define printed circuit boards and standard parts libraries. Multi-page drawing windows and menu-driven command structure simplify creation of complex mechanical drawings. Automatically dimension your fabrication drawings. Output release-quality documentation.

Versatec and Expert are trademarks of Versatec, Inc.
Xerox is a trademark of Xerox Corporation.

EXPERT FOR PCB DESIGN.



VERSATEC
A XEROX COMPANY

VERSATEC
ENGINEERING INFORMATION SYSTEMS
2710 Walsh Avenue
Santa Clara, California 95051
Telephone: (408) 988-2800
TWX: 910-338-0243
Telex: 334421

Versatec Expert gives you four integrated software modules for printed circuit board design. Design, simulate, place and route, and complete all manufacturing documentation at any single Expert workstation or share applications throughout the Expert Ethernet network.

Ethernet connects Expert workstations to create an integrated engineering automation system that can support over a thousand users. High-speed communications enables sharing of a common database, centralized files, and peripherals.

The Expert PCB design workstation is backed by nationwide Xerox service. Low entry cost, high performance, and attractive leasing terms make Expert the easy choice for PCB design.

Circle our readers' service number for more information. Better yet, call your closest regional demonstration center for your own hands-on demo.

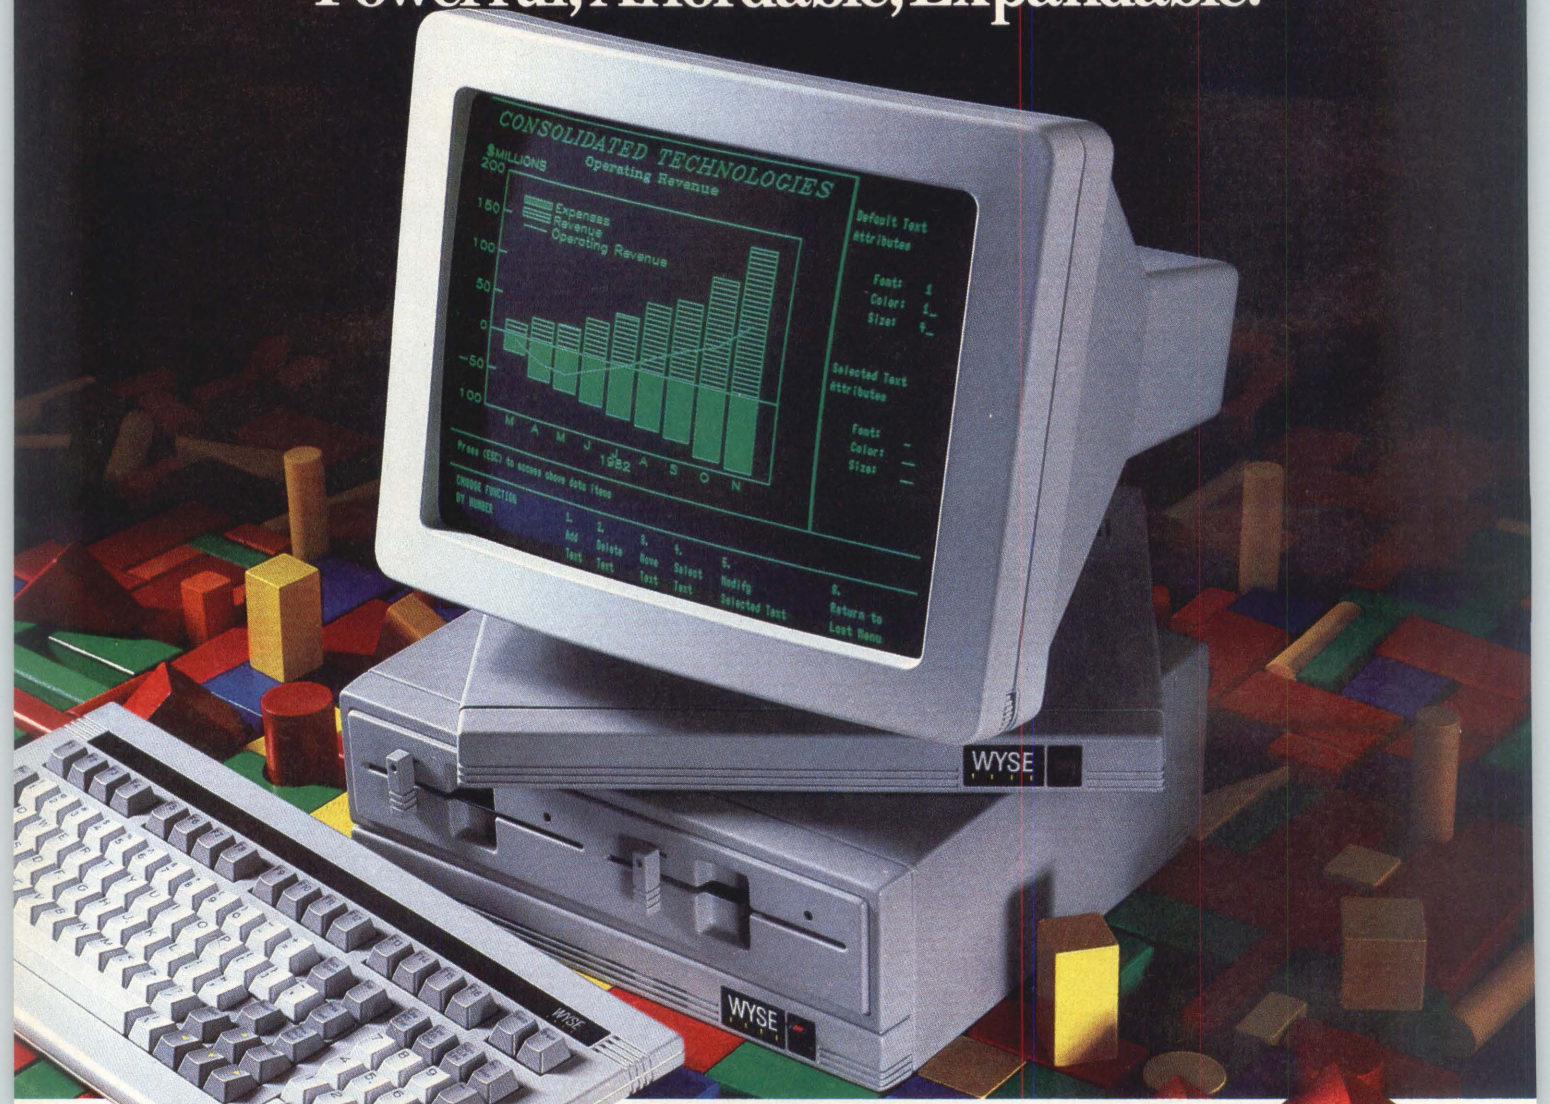
REGIONAL DEMONSTRATION CENTERS

Atlanta, GA	404-992-3198
Boston, MA	617-229-6747
Chicago, IL	312-885-2757
Dallas, TX	214-620-8694
Indianapolis, IN (DSI)	317-299-9547
Maitland, FL	305-660-2126
Newport Beach, CA	714-851-8005
San Diego, CA	619-452-5611
Santa Clara, CA	408-244-5581
Wayne, PA	215-293-0920

EXPERT. THE EASY CHOICE FOR PCB DESIGN.

CIRCLE 30

The New WY1000 Microcomputer
BUILDING BLOCKS
Powerful, Affordable, Expandable.



The WY1000 stacks up to be a lot of machine from a few simple pieces. By adding the WY1000 microcomputer to the good-looking, ergonomic WY50 display terminal, we created the most exciting concept in desktop workstations on the market today.

We also added sophisticated high resolution graphics, suitable for the most demanding applications.

Plus, we added color capability, when used with our color terminal.

And on top of that, we added a Winchester Disk Drive option providing an additional 10 megabytes of storage.

FEATURES:

- **80186 16 Bit 8 MHz Processor**
- **128KB to 768KB RAM Memory**
- **Two Floppy Disk Drives (725 KB)**
- **Optional 10 MB Winchester Drive**
- **RS232 & RS422 Serial Ports**
- **Optional Graphics/Color Graphics**
- **Networking Capability**
- **CP/M™, MS-DOS™ Compatible**
- **Priced from only $1995**

Best of all, we priced the WY1000 from only $1995. It all adds up to a system builder's dream.

For a complete brochure on the WY1000 contact Wyse Technology toll free at 800/421-1058.

WYSE

Make the Wyse Decision.

WYSE TECHNOLOGY *3040 N. First St., San Jose, CA 95134, 408/946-3075, TLX 910-338-2251, Outside CA call toll-free, 800/421-1058, in So. CA 213/340-2013.*

CP/M is a registered trademark of Digital Research Corporation. MS/DOS is a registered trademark of Microsoft Corporation.

Multitasking Unix look-alike supports 11 IBM XT users

Many computer designers believe that the last thing they need is yet another Unix look-alike to add to the bewildering array of similar operating systems. However, an alternative now exists that makes a multi-user Unix available on the IBM PC XT and its clones. This look-alike also offers runtime or production operating system versions at a greatly reduced price and furnishes the tools for easy program development of multitasking, multi-user Unix Version 7 compatible software.

Network Consulting, Inc (Discovery Pk, Suite 110, 3700 Gilmore, Burnaby, British Columbia, Canada V5G 4M1), has taken the Coherent operating system, a two- to three-user multitasking Unix look-alike (see *Computer Design's* Special Report on Operating Systems, July 1984, p 153), and adapted it to a multi-user environment on the industry standard PC XT. Using the Unix Version 7—which is compatible with the NCI Coherent—and hardware added to the IBM PC XT, up to 11 users can simultaneously operate the PC XT. The system can handle many vertical applications, including order entry, record keeping, and other office activities (see Panel, "A practical system").

To design its operating system, NCI, under license, made major changes in the Mark Williams Co's (Chicago, Ill) two- to three-user Coherent operating system. Coherent itself, was developed independently from the original AT&T Bell Labs Unix Version 7, but was designed to be a look-alike. In particular, says NCI marketing engineer Dale A. Thomlison, parts of the original Coherent kernel have been rewritten for faster multitasking operations. For example, new serial line drivers, disk driver DMA interrupts, and schedulers are all the result of the NCI design.

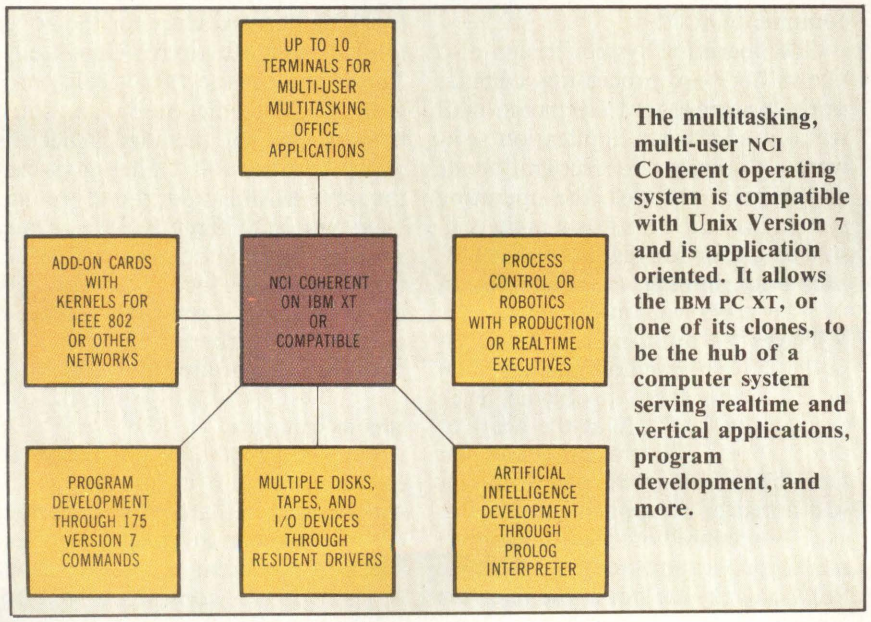
There are a variety of other modifications implemented with the goals of making the NCI operating system Unix-compatible and optimizing the system for designers and software developers. According to NCI vice president of research and development Angus E. Telfer, the 165 to 175 Unix Version 7 commands are geared toward allowing software developers to create Unix-compatible applications to run with either a runtime or production operating system version.

In order to ensure faster execution of Unix Version 7, Telfer rewrote the Coherent mathematics library in assembler and reworked some features of the kernel code, including the scheduler and the C library. The CAT command (to concatenate and print files at the user's terminal), for example, now executes two to three times faster. According to Telfer, the mathematics library in the proposed IEEE floating point format, is the fastest floating point routine—proven by its benchmarks—that he has seen thus far.

NCI Coherent incorporated such Unix system memory routines as memcopy, memchr, memcmp, memcpy, and memset, which ensure memory locks and shared memory. These make Unix System V compatible with the multitasking NCI Coherent kernels that Telfer wrote for various IBM XT add-on cards. The add-on cards such as an Arcnet card and the 8088-based Persyst smart serial interface card. Computer designers opting for this hardware can enjoy, for example, RS-232-C based networking with hardware handshaking and flow control. By design, realtime interrupts go from the Persyst card by message passing to the XT's operating system kernel.

Telfer chose to pack the operating system with software development tools. For example, he has included a multipass assembler intended for use as a target by the NCI Coherent C compiler. The assembler is available only for small assembly routines and supports Intel's iAPX-86 assembler. The operating system also sports a Prolog interpreter written in C for artificial intelligence (AI) applications. With this C interpreter—similar to the University of Edinburgh's DECsystem 10 Prolog developed by Claude Sammut at the University of New South Wales, Australia—the system integrator can enjoy a Unix-based AI language for PC XT-based AI research and development projects. With additional work, the system can accommodate a multi-user Unix interface with the Lisp language through one of the workstation

(continued on page 54)



The multitasking, multi-user NCI Coherent operating system is compatible with Unix Version 7 and is application oriented. It allows the IBM PC XT, or one of its clones, to be the hub of a computer system serving realtime and vertical applications, program development, and more.

Unix look-alike

(continued from page 53)

A practical system

Some industry observers feel that the single-user Unix developed by Interactive Systems (Santa Monica, Calif) for IBM's PC XT is just a toy—largely because of its speed limitations for I/O intensive applications. Well aware of this feeling within the computer design community, NCI has taken pains to explain just what it means by multi-user operation and what kind of performance can be expected from its operating system (see the Table).

Although Dale Thomlison of NCI feels "there are lies, damn lies, and benchmarks," he contends that the NCI Coherent's benchmarks prove the operating system can convert the PC XT into a practical, multi-user system for certain vertical applications. The operating system is a capable single- to two-user system when the XT is to do compilation intensive development work, according to Thomlison.

It is good for three to five users when performing tasks such as editing and word processing, Thomlison adds. Moreover, it can handle 10 or 11 users if a single application is accessed or if it is used as a telephony controller.

Benchmark Comparisons

Machine	Disk test (read-write)	Piper	Procedure Calls	Sieve Test	System Calls
11/780	10.2/23	2.6/6	31.9/32	4.8/5	1.8/2
11/750	14.5/24	4.2/8	52.5/53	8.7/9	3.0/3
Codata	16.0/73	4.4/11	43.7/44	7.5/8	3.5/5
11/34	127.0/310	15.7/111	97.0/98	17.2/18	8.6/9
* IBM/XT	65.1/117	40.6/54.4	93.2/94	18.1/18	12.5/13

* NCI COHERENT
[times are in seconds (CPU time)/(elapsed time)]
(Not latest version)

computers that handles applications. These workstations are offered by LISP Machines, Inc (Culver City, Calif) and Symbolics, Inc (Chatsworth, Calif).

Put it to use

All these features have been incorporated with an eye toward designing an application-oriented operating system that presents more than just an elaborate monitor program. For example, the PC XT can be used as a database-oriented machine such as a data multiplexer. It has such built-in networking features as login, password, encryption, and file protection. The XT can also serve as a protocol converter, allowing up to 50 different kinds of terminals on a multi-user system to communicate using the Berkeley termcap data base. Finally, it can monitor a private branch exchange and log telephone calls.

The operating system handles process control and robotics functions because it has a realtime executive that can be burned into PROM. Here, Telfer said, real time might mean 5 s in a process control application. The executive is geared to operate on intelligent add-on cards that serve as IBM PC XT I/O processors. The Persyst card, for one, can run four I/O lines at 19.2-Kbyte rates. This operation can use RAM-disk support.

Other Coherent modifications of interest to the computer system designer include a screen editor and a source-code control system. The former allows easier application development while the latter allows developers to keep track of different versions of the code they have written. Further capabilities for software developers include a lexical analyzer (LEX) and a compiler-compiler (YACC).

The operating system design also allows for word processing, calculations, printing, and electronic mail without additional application software. Moreover, in recognition of marketplace realities, the operating system uses the PC XT as a main console emulator for reading MS-DOS 1.1 files and running 1.1 programs. A DOS 2.0 version is on the way.

The PC XT multi-user capability is said to have no memory management or execution speed problems, even though the Intel 8088 at the heart of the machine has no provision for Unix-like memory management. It is just a matter of applications, Telfer said. He pointed out that the vertical market applications for the 2.5-Mbyte NCI Coherent are not compile or I/O intensive, adding that customers of the typical application program for dental and insurance companies, and similar services with specific needs have had no problems with it. "I wouldn't expect a lot of compiles to go on in applications," he said, noting that program development in an environment requiring heavy compiling is a one- or two-user operation.

The NCI operating system uses 2.5 Mbytes compared to 5 Mbytes for the Interactive Systems' single-user, multitasking licensed Unix for the PC XT (and about 8 Mbytes for Unix itself). This is because the TROFF utility program, the Unix dictionary, and some macros are not included. Space is saved because NCI Coherent is not the same as Unix and, therefore, has been written differently. Use of even shorter run time or smaller production kernels—which NCI will help develop for computer system designers—further minimizes the code size, according to Telfer.

Lots of hardware

One of the major chores of an operating system is to allow for a wide variety of hardware to be used on the computer on which it operates. The NCI Coherent is no exception.

(continued on page 56)

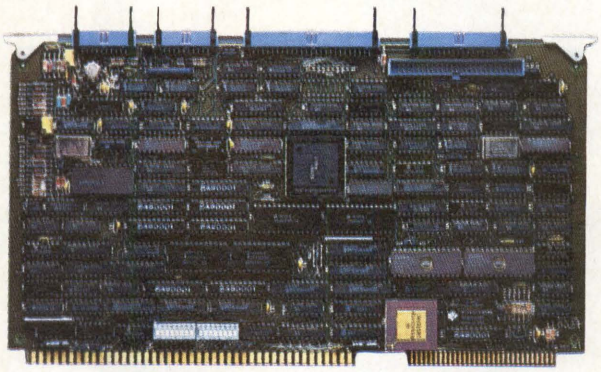
"With the Interphase Storager,™ I can make a 5¼" hard disk perform like an 8" disk."



Frank Emser
Manager Hardware Development
Paradyne Corporation

The Interphase Storager Multibus™ controller can give a 5¼" Winchester disk capabilities never before possible. Storager not only gets more performance from existing ST506 drives, but also supports the new ESDI and ST412HP interfaces for more power and capacity than ever before. And since Storager can control two Winchester disks, four ¼" tapes (QIC-02), and two 3½", 5¼" or 8" floppies, the same controller can be used for every storage need.

Storager features 1:1 interleave, with concurrent disk and tape transfers and simultaneous disk and bus transfers for speed and high performance. And Storager's unique "virtual buffer" architecture with UNIX™-optimized intelligent caching can reduce or eliminate disk rotational latency and overcome data overrun/underrun problems of FIFO-based controllers. Plus, for the very first time on a controller, Storager has an *on-board* 68000 CPU.



The Storager controller is the latest product in Interphase's line of high-performance Multibus controllers. Interphase also offers Multibus controllers for SMD disks, local area networks and video monitors. Plus powerful disk controllers for the IBM® PC. They're all backed by a great customer support team that works full time with Interphase customers to assure that our products work the way they should — in the system.

Find out how Storager can make a 5¼" disk perform like an 8" disk. Call Interphase today at (214) 350-9000.

Storager is a trademark of Interphase Corporation
Multibus is a trademark of Intel Corporation
UNIX is a trademark of Bell Laboratories
IBM is a registered trademark of International Business Machines

INTERPHASE corporation

2925 Merrell Rd. • Dallas, TX 75229

CIRCLE 32

Unix look-alike

(continued from page 54)

For example, NCI Coherent supports the control system, an 8-line serial port card for the PC XT. If two regular IBM serial lines are added, 11 users can work. Both X-ON/X-OFF and clear to send (CTS) handshaking with 5, 6, 7, or 8 data bits, with or without parity, may be used in either an interrupt or polling mode. Other hardware supported by drivers built into the operating system include IEEE 488 interfaces to talk with any controller satisfying this protocol, some dozen hard disks, most PC XT compatibles, standard storage module device (SMD) gigabyte disks, nine-track tape reads or writes, and the previously mentioned smart communications cards. The operating system allows over 50 different kinds of terminals to communicate with it through use of the Berkeley termcap data base.

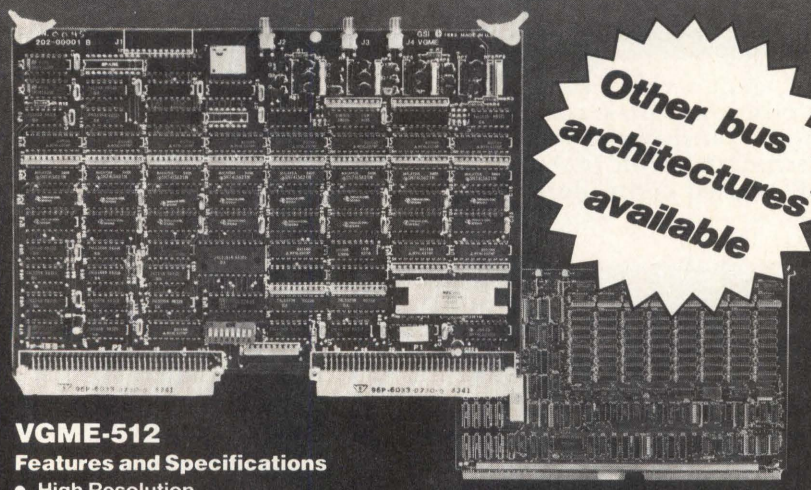
The licensing fee for NCI Coherent is included in the purchase price. The total price is far less than what is charged by licensed Unix look-alikes because AT&T does not receive a share of NCI Coherent's profits. Note that the fee includes the full multiuser facility—there is no extra for additional users. Even less expensive (less than $100 in quantity) are the runtime and production versions.

The NCI Coherent operating system allows a computer network system to be IBM PC XT-based with up to 10 remote terminals at a price competitive with that of a dedicated multi-user Unix system, according to Thomlison. Depending on the hardware configuration used, the NCI Coherent operating system could be more cost effective than multiuser Unix systems. The PC XT with NCI Coherent is 1/50 the cost of the DEC VAX traditionally used for Unix-compatible package development, has about 1/10 the power, and needs no support.

Plans for NCI Coherent may include moving it into the Unix marketplace with its large array of add-on hardware device support for the PC XT and its clones. The firm is also working to develop NCI Coherent-based application software. And, as a specialist in fast p-System software, NCI is working to port p-System applications to its new multitasking, multi-user operating environment.

—Harvey J. Hindin,
Special Features Editor

VME
Graphics Controller



Other bus architectures available

VGME-512
Features and Specifications

- **High Resolution**
  512 x 512 Pixels
  Dot Addressable via Cartesian Coordinates
  4 Bits-Per-Pixel
- **Color/Intensity**
  Eight Colors with Dual Intensity
  Blink at 1Hz to Black or Reduced Intensity
- **High Performance**
  800,000 Pixel-Per-Second Vector Generator
  All Planes Drawn Simultaneously
  60Hz Non-Interlaced Refresh Rate
  Line by Line Vertical Scroll
  Multiple Cursors of Arbitrary Shape
- **GPS Subroutine Package**
  Points, Lines, Circles, Arcs, Rectangles and Polygons
  Line and Area Texturing
  Alphanumerics—96 Characters—2 Sizes
  Picture Storage & Retrieval

GSi Graphic Strategies, Inc.
1445 Koll Circle, Suite 107
San Jose, CA 95112
(408) 294-1300

SYSTEM TECHNOLOGY
(continued on page 60)



How we created an MDS that looks like a million, performs like 100 grand and costs under $20,000.

Introducing Emulogic's ECL-3200 universal development system.

It wasn't simple. And it took a lot more than seven days.

But when you start with the world's most advanced emulator, you're already way ahead. The ECL-3211 universal emulator gave our new MDS a repertoire of features you can't find in any comparable system. Like full-speed, no-wait-state emulation. Eight real-time, multi-function breakpoints for every chip. Real-time trace decode. And full-speed, full-range memory mapping with single-word resolution.

And then DEC™ made our job a little easier, too. We built their powerful new desktop computer–the Professional 300–right into the system. With RT11XM,™ 11/23 with MMU, 22-bit addressing, 512KB memory, 800KB dual diskette and optional 5 or 10MB Winchesters.

We didn't skimp on looks, either. We wrapped the ECL-3200 in a clean, compact package that, frankly, we call beautiful. But you tell us what you think.

You might conclude that, with all we've put on the inside, the price tag on the outside would be just as impressive. Not so! We brought the ECL-3200 in for 25% less than expected. So now you can buy a full-function development system for under $20,000–25% less than you'd expect. Including an 8-bit emulation package for the Z80,® 6502, 6809 or another popular chip. And you can also use it with all 16-bit emulation packages and our new high-level software development tools like C and Pascal cross-compilers and our just-introduced SLICE™ Symbolic Debugger.

So you see, creating our miracle MDS wasn't all that difficult. But we don't recommend anyone else try.

For detailed information, call 800-435-5001 or 617-329-1031 in Massachusetts. Or write Emulogic,® Three Technology Way, Norwood, MA 02062.

Professional 300, RT11XM and DEC are trademarks of the Digital Equipment Corporation. SLICE is a trademark, and Emulogic is a registered trademark, of Emulogic, Inc. Z80 is a registered trademark of Zilog, Inc.

EMULOGIC®
MICROPROCESSOR
DEVELOPMENT SYSTEMS

European Distributors: Austria: Walter Rekirsch, (43 222) 235555; Denmark: Instrutek, (45 5) 611100; France: YREL, (33 3) 9568142; Sweden: Aktiv Elektronik AB, (46 8) 7390045; Switzerland: Instrumatic AG, (41 1) 7241410; United Kingdom: MSS, (44 494) 41661; West Germany: Instrumatic Electronic GmbH, (49 89) 852063.

CIRCLE 34

© 1984 Quantum Corporation

We'v
increa
our le
to 5¼

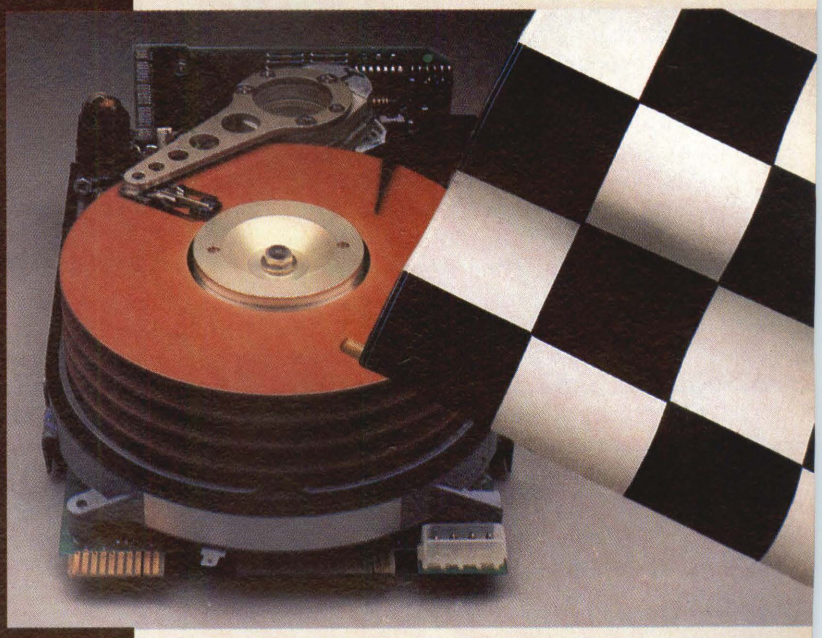
e sed ad.

The clear lead we established with our 8" medium-capacity Winchesters has now been extended. Our 5¼" drives are out front now, too. With both 8" and 5¼" drives leading the field, we now make more medium-capacity disk drives than anybody.

If there's a lot riding on your disk drive decision, you may want to consider our record as a consistent winner.

5¼" drives, from 20 to 40 megabytes. 8" drives, from 10 to 85 megabytes.

Quantum Corporation, 1804 McCarthy Boulevard, Milpitas, CA 95035, (408) 262-1100, TWX 910-338-2203. Eastern Regional Sales Office: Salem, NH (603) 893-2672. Western Regional Sales Office: Santa Clara, CA (408) 980-8555. European Sales Office: Frankfurt, West Germany 611-666-6167.



QUANTUM

CIRCLE 35

Flat panels approach monochrome CRT specifications

In the efforts to increase size and resolution of today's flat-panel displays, active matrix technology has made recent breakthroughs. These developments promise an avenue to flat panels that can rival CRTs, at least for monochrome displays. The active matrix approach places an active element, in the form of an electronic switch, at each pixel. The most common such switches are thin-film transistors or metal-insulator-metal nonlinear resistors. The aim is to enable update of the display at video rates and with CRT resolution.

Since good gray scale is vital to pocket LCD television, several other active matrix technologies are being investigated. In addition to thin-film transistors (TFTs), and metal-insulator-metal (MIM) resistors, diode rings are candidates for active matrix applications. MIMs are nonlinear devices that can be driven at video rates and yield a gray scale.

The diode ring approach appears to offer a very small variation of threshold voltage, is simple to reproduce, and can be used to build displays of up to 5000 lines. Displays implementing gray scale have actually been produced using the diode ring as the active element.

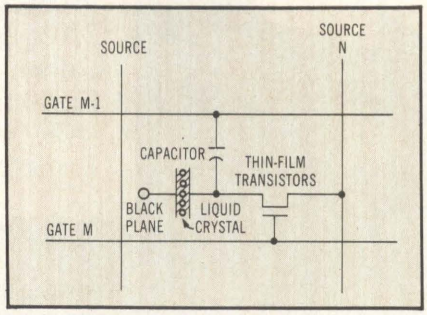
More elusive is the quest for the ultimate flat-panel display—the one that will offer the same resolution and color as a color CRT. Significant strides in that direction have been made in Japan. These were discussed at the recent conference of the Society for Information Display (SID) in San Francisco by the Suwa Seikosha Co, Ltd (Suwa, Japan) and Hosiden Electronics Co, Ltd (Osaka, Japan).

The Seikosha display, a 4¼-in. (10.8-cm) panel with 480 x 480 pixels, features TFT arrays. The display uses back lighting through the LCD array, which acts as a light valve for a color filter. In the Hosiden display, microcolor layers of red, green, and blue are aligned with corresponding TFT pixels, yielding a still rather grainy color display.

The liquid crystal medium

While various methods of producing active matrix black and white displays were also presented at the conference, a flat-panel display incorporating TFT with a liquid crystal medium has been announced by Panelvision Corp (265 Kappa Dr, Pittsburgh, PA 15238). Although Panelvision designers choose liquid crystal media, the TFT active matrix approach lends itself to all forms of display media currently used.

In fact, the real breakthrough of active matrix technology lies not so much in the display material as in its ability to address and drive a large number of display elements. With LCDs, each pixel can operate at 100 percent duty cycle, thus achieving very sharp contrast and viewing angle in comparison to time division multiplexed LCDs. The TFT approach also lends itself to eventual incorporation of gray scales.



The actual transistor takes up a small part of the pixel area in this one-pixel elementary cell for the active matrix display. Vacuum deposited on glass, the cell can be replicated according to the size limits of the vacuum equipment.

Multiplexed dot-matrix LCDs have an inherent limitation in that, as the number of multiplexed lines increases, the differences between the on and off voltages for each line of pixels decreases and the contrast between pixels degrades. One way to combat this is to increase the drive voltages as the number of lines increases, but this too has obvious limitations. In addition, since multiplexed LCDs form passive arrays, the nature of the material determines threshold voltage, response, and switching speed.

The unit announced by Panelvision features an active area of about 4 x 3 in. (10 x 8 cm) and displays 512 chars using a 5- x 7-pixel matrix. The pixels total 24,576. In addition to the 512-char product, Panelvision is currently engineering an 80-char, 25-line display that will be a 400-line resolution bit-mapped matrix. As the displayed information approaches that of a standard 24- x 80-char CRT, the number of pixels that must be directly addressed approaches 250,000.

Matrix resembles a fully decoded RAM

Panelvision has solved this by creating what is, in effect, a 5¼- x 4-in. (13- x 10-cm) IC. The TFTs are vacuum-deposited on a glass substrate that takes up an arbitrarily large area dependent only upon the size of the vacuum equipment. The LCD material is injected and the panel sealed. The same kind of TFTs can be used to produce row and column drivers and shift registers, thus integrating peripheral circuitry onto the display panel. This reduces the number of connections needed for external circuitry to about 10. The entire matrix is analogous to a large, fully decoded RAM. Since the full duty cycle is available at each pixel, there is room to introduce levels of gray scale in later products.

Panelvision designers emphasize that active matrix technology is independent of the display material and can theoretically be used for plasma, electroluminescent or LED displays. These materials, however, have their own inherent limitations in terms of resolution, power consumption, and threshold.

For example, electroluminescent displays have a sharp on/off threshold. This allows multiplexing of more horizontal lines, but it does not allow for gray scale control of the pixels on the lines. On the other hand, the threshold characteristics of liquid crystal allow sufficient range because the material appears to modulate the light gradually as voltage is increased. As other materials exhibiting the desired characteristics are developed, the active matrix addressing method should make use of them.

—Tom Williams,
West Coast Managing Editor

CONNECTORS, WITH OUR ADVANTAGE.

The Contact Advantage:

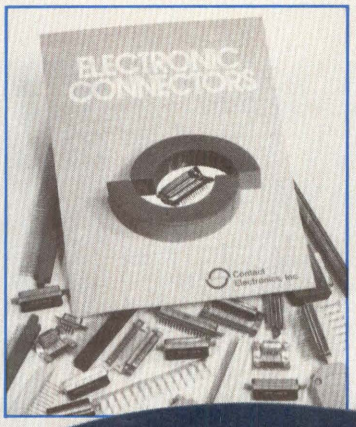
- *second source interchangeability* gives you an economical alternative in meeting your electronic connector needs
- *major inventory availability* means shorter lead times in filling your orders
- *local distributor network* for convenient localized service whenever, wherever needed
- *rush hotline service and guaranteed delivery* (call us for details)
- *connector quality you can count on* —quality *built* in, not just "inspected" in; UL recognized.

D' Subminiature Connectors: 9-50 Position Solder Cup, Solder Pin, Wrap Post, Crimp (Insulators & Contacts); 9-37 Position PC Mount Right Angle. Conform to MIL-C-24308.

IDC Connectors: 9-50 Positions; quick economical mass termination with round conductor flat ribbon cable.

Eurocard-DIN Connectors: Solder Pins, Solder Loops, Angled Solder Pins, 2-3-Level Wrap Post, Angled Wrap Post. Conform to DIN 41612 & 41617.

Cardedge Connectors: Solder Pin, Eyes or Solderless Wrap Terminations. Conform to MIL-C-21097.



For more information on our electronic connectors, hoods, mounting hardware and other products offering money-saving, short-lead-time, second source interchangeability and other benefits with our advantage, ask for our catalog. Just write or call us toll-free:

800-221-1451



contact®

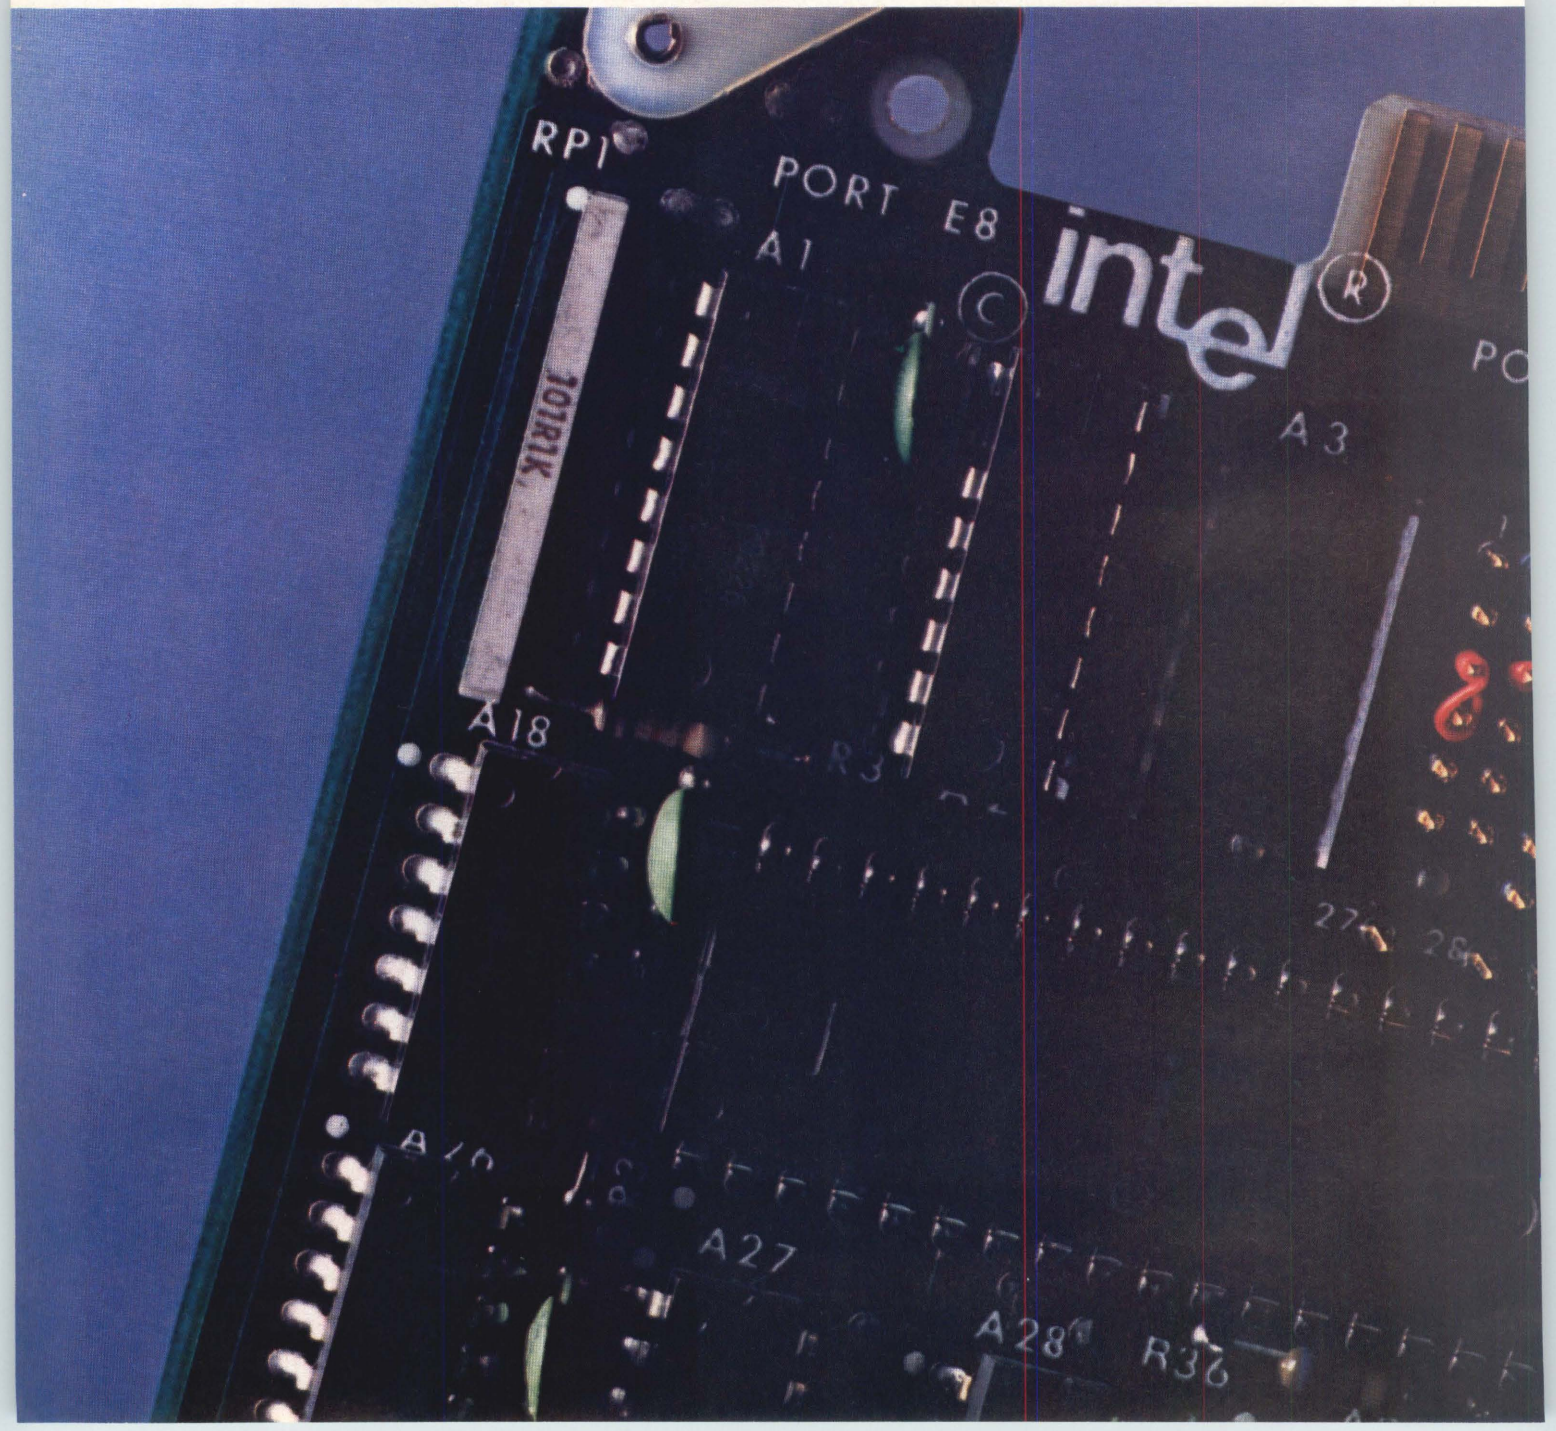
Contact Electronics, Inc.

30 Plymouth Street, Fairfield, NJ 07006 • Phone 201-575-7660 • Telex 130-216
Western Regional Sales Center: 1900 Wyatt Drive #10, Santa Clara, CA 95050 • Phone 408-988-5400

CIRCLE 36

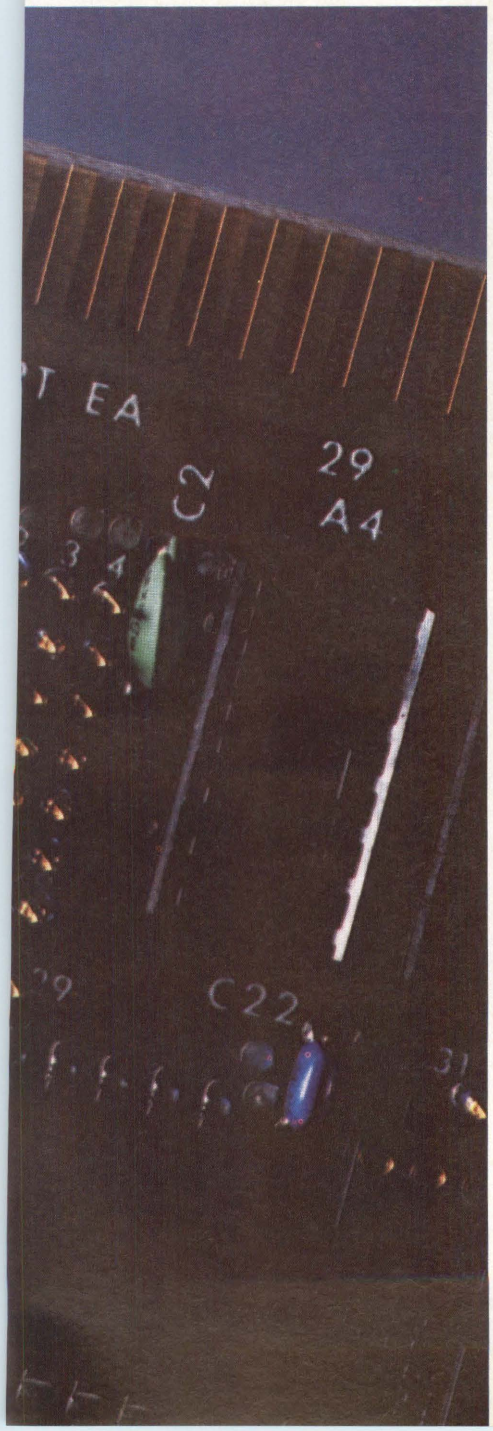
CEI-5

THE BEST PLACE TO BOARD THE BUS.



Of course, we're talking about the MULTIBUS® architecture. The most widely accepted, best supported bus architecture in the world.

But since you've already made the wise decision to use MULTIBUS as your road to riches, we're now going to show you the best place to start your journey.



Right here. Because as you might suspect, we know just about all there is to know about MULTIBUS. After all, we invented it. And we have the largest selection of MULTIBUS products to be found. Anywhere.

For starters, there are our boards. We've produced over a million since 1976. Which should give you an indication of the experience we have in manufacturing quality boards. Quality that's not easily achieved.

Our boards are scrutinized at every level of integration. From chip through system level. And each board is tested to make sure it does its job in every one of its configurations. We even pull finished product from our warehouse for retesting and inspection. Just to make certain nothing slips by.

You can count on the road being smooth for a long time to come, too. Reliability is assured from start to finish. We continually monitor each board throughout its entire production life. Even going so far as lot sampling older boards and putting them through 5000-hour life test studies.

Then there's the selection. A big one.

Over 115 different board-level products. Supporting 15 CPUs (from our 8080 to 80286). So you won't get stuck into one or even two design approaches. A tranquil thought when you consider how fast things change in the marketplace.

What's more, you can count on getting everything you need. From peripheral controllers to memory expansion, data communications and graphic boards. Plus card cages and software (including iRMX 86™, the most popular real-time operating system in the MULTIBUS world).

One more fact about our MULTIBUS products. You can get fast service anywhere along the road. We've got one of the largest field engineering forces in the world. Over 600 factory-direct people. Not to mention localized support from our worldwide network of trained distributors.

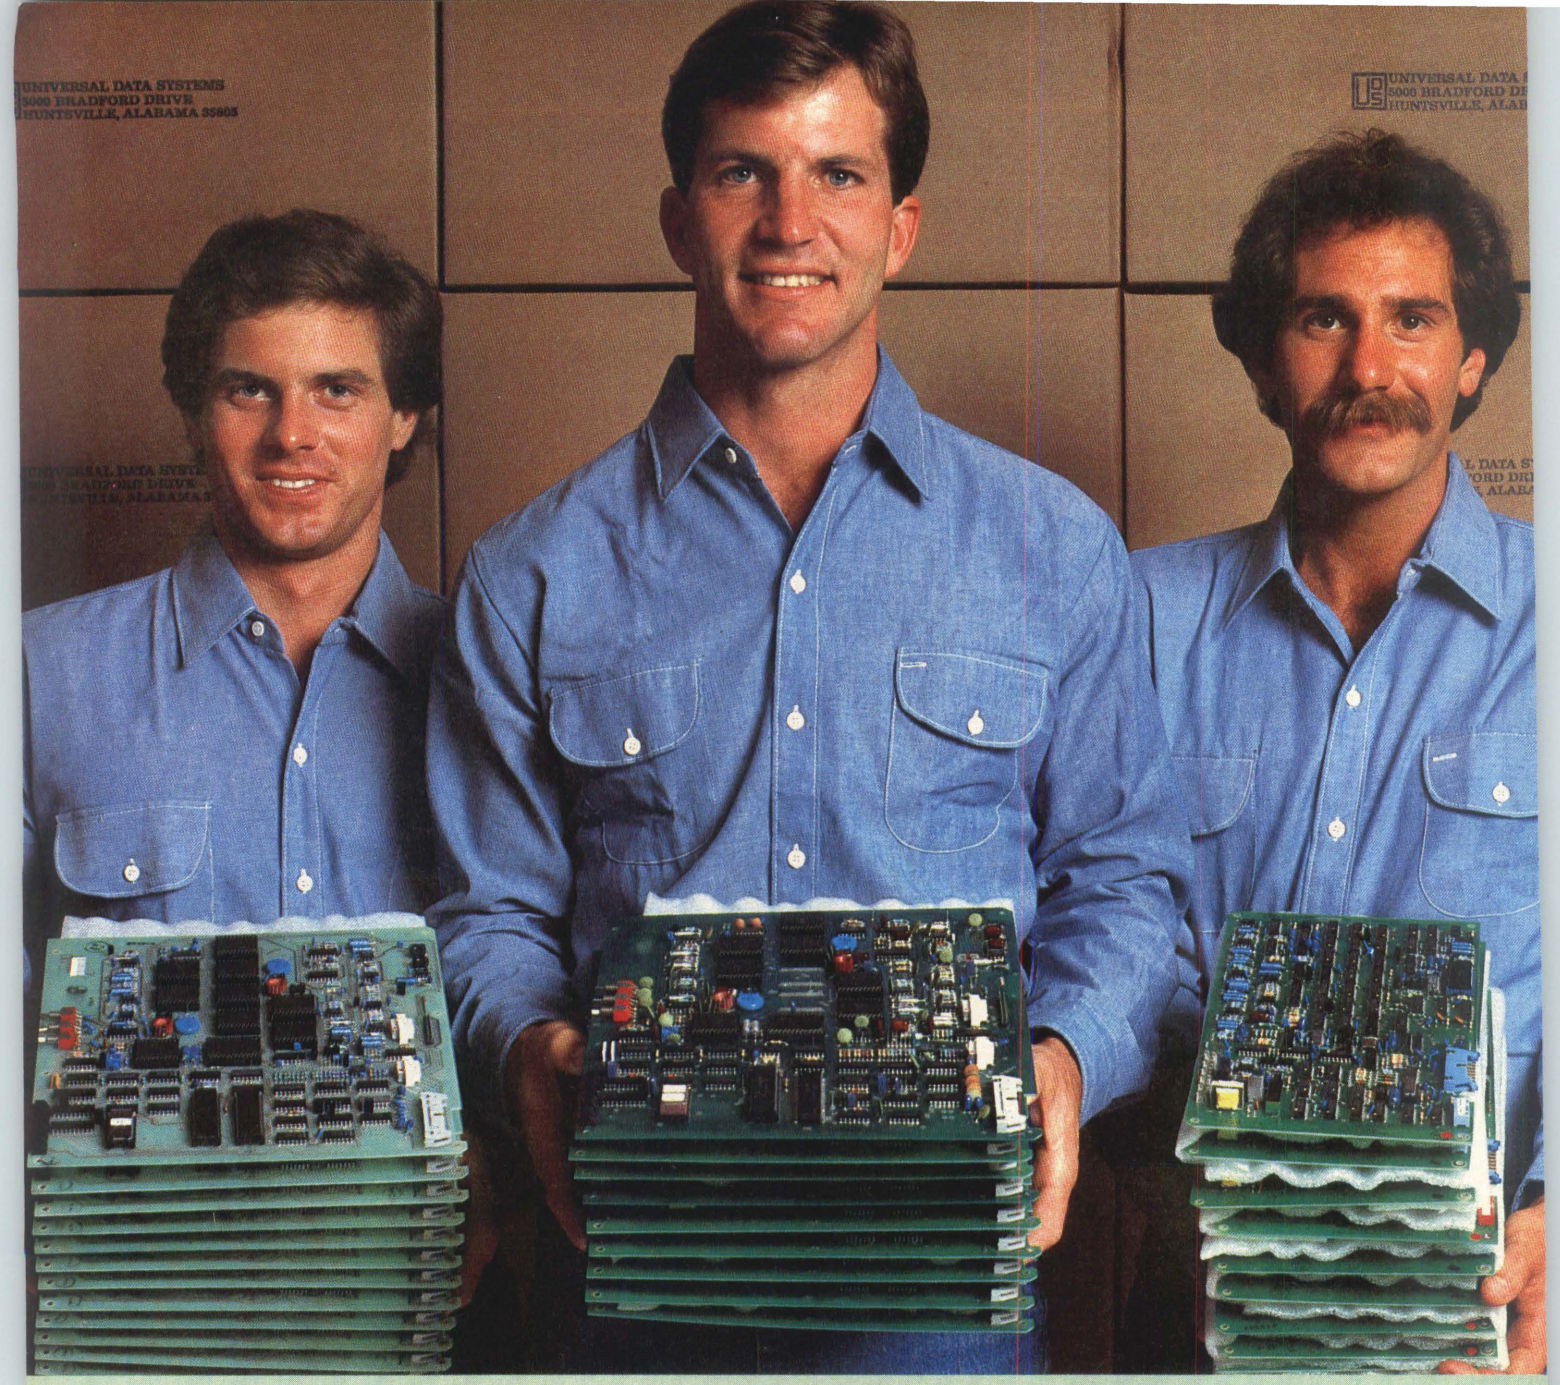
So don't miss the bus. Travel safely. And arrive sooner. Call toll-free (800) 538-1876. In California, (800) 672-1833. Or write Intel, Lit. Dept. C-20, 3065 Bowers Ave., Santa Clara, CA 95051.

We'll be happy to welcome you aboard.

intel®

© 1984 Intel Corporation

CIRCLE 37



Chairmen of the boards!

The UDS shipping crew has directed a lot of boards! Modem boards for a multitude of OEM customers are on their daily agenda.

Upstream from our chairmen of the boards is a highly automated, high-volume modem manufacturing facility, and a staff of industry leaders in OEM modem design and economical board layout.

UDS modem boards are available in 103, 201, 202, 208 and 212A Bell-compatible versions, as well as a 9600 bps modem that's operable over either dial-up or dedicated lines. All our dial-up versions are FCC certified for direct connection to the switched network.

In low-cost standard configurations or space-saving custom layouts, UDS OEM modems set the highest standards for reliability and cost effectiveness. Prove it to yourself with your next modem buy. Contact Universal Data Systems, 5000 Bradford Drive, Huntsville, AL 35805. Telephone 205/837-8100; TWX 810-726-2100.

Universal Data Systems

MOTOROLA INC.
Information Systems Group

UDS modems are offered nationally by leading distributors. Call the nearest UDS office for distributor listings in your area.
DISTRICT OFFICES: **Atlanta,** GA, 404/998-2715 • **Bellevue,** WA 206/455-4429 • **Blue Bell,** PA 215/643-2336 • **Boston,** MA, 617/875-8868 • **Columbus,** OH, 614/895-3025 • **Englewood,** CO, 303/694-6043 • **Glenview,** IL, 312/998-8180 • **Houston,** TX, 713/988-5506 • **Huntsville,** AL, 205/837-8100 • **Mountain View,** CA, 415/969-3323 • **Old Bridge,** NJ, 201/251-9090 • **Richardson,** TX, 214/680-0002 • **Silver Spring,** MD, 301/587-0166 • **Tampa,** FL, 813/684-0615 • **Tustin,** CA, 714/669-8001

Created by Dayner/Hall, Inc., Winter Park, Florida

CIRCLE 38

Midcon/84 High Technology Electronics Exhibition and Convention

Mini/Micro Southwest Computer Conference and Exhibition

Dallas Convention Center, Dallas, Texas **September 11 to 13, 1984**

This year, Dallas, Texas will be the location for the jointly held Midcon and Mini/Micro Southwest Conferences and Exhibitions, on September 11 through 13. Both will be housed in the Dallas Convention Center. In addition to separate Professional Programs for Midcon and Mini/Micro, technical products will be on display at more than 900 booths in the Convention Center.

This year's speaker at the Keynote Luncheon will be Admiral B. R. "Bobby" Inman, USN (retired), who is currently president, chief executive officer and chairman of the board of directors for Microelectronics and Computer Technology Corp (MCC) of Austin, Texas. The Keynote Luncheon is a separate function with a $15 fee and will be held on Tuesday, Sept 11.

Midcon has a much larger technical program than does Mini/Micro Midwest—20 sessions compared to nine—but only half of Midcon's agenda is aimed at the system designer. On the other hand, all of Mini/Micro's sessions touch on some phase of digital technology. Midcon sessions that are technology oriented include such diverse subjects as speech technology, surface-mounted technology, programmable logic, communication network design, testing of programmable logic devices, the integrated services digital network, nonvolatile memory, emi/rfi regulations, digital signal processing chips, and computer aided engineering.

Mini/Micro sessions cover integral modem design, serial protocols, CMOS single-chip microprocessors, VLSI bit-mapped graphics controllers, multiprocessor systems, system-increasing system throughput, microprocessor development, 16-bit development tools, and high performance Multibus systems. Some are design oriented, others are application oriented, but all involve the system designer.

Midcon/84 technical sessions

At Midcon, the Tuesday early morning and late afternoon "Premier Sessions" are probably the only ones on that day that will be of primary interest to the design engineer. However, because these sessions are considered to be of importance, both will be repeated on Wednesday and Thursday for registrants who cannot attend the sessions on Tuesday.

The first of these Premier Sessions, from 9:30 to 11:30, acknowledges the importance of speech as the ultimate interface between humans and machines. Speakers will discuss system- and board-level advances made in speech technology—covering both recognition and generation. The session organizer and speakers are from semiconductor innovators such as Motorola, Texas Instruments, and American Microsystems. As might be expected, discussions will be divided between prototype and production devices, but emphasis will be on widespread application of speech technology—advancing speech technology from novelty status to accepted, everyday usage.

Recent proliferation of available components, from SSI to VLSI circuits, has increased interest in surface-mounted technology. Considerations and application of this technology in the design and manufacture of printed circuit boards will, therefore, be covered in the late afternoon Premier Session. Emphasis of the papers will be on the design considerations, particularly on ramifications for the future.

Wednesday will offer a much wider range of sessions oriented to the design engineer. Among these are the application of programmable logic devices (PLDs) and communication network design in the morning; and testing of PLDs, the integrated services digital network, nonvolatile memories, and emi/rfi regulations in the afternoon. These sessions are in addition to the repeated Premier Sessions mentioned above.

Basis for the sessions on PLDs is the prediction that their use will increase by 10 times in the next few years. The early morning session will explain just what PLDs are, what their benefits are for designers, and where they fit into the design of future subsystems. Emphasis will be on software and programming as well as designing-in testability. In the early afternoon session, speakers will discuss

(continued on page 67)

ARC's new INTERVIEW® COMSTATE™ data comm protocol analyzer family.

The ultimate in programming power and simplicity.

ARC's new INTERVIEW COMSTATE data protocol analyzers are the ultimate in high level programming power yet simple to program. Their new programming technique is so logical to use that after only a one hour introduction, you will be able to develop effective and previously difficult to program tests. So why spend days reading manuals and learning how to program other testers? Evaluate the COMSTATE family now.

The COMSTATE programming technique parallels the latest protocol development tool—State Diagrams. Because of this, our COMSTATE analyzers are easier to program and they emulate data communications protocols more precisely than any other analyzer/emulator.

COMSTATE Programming. A powerful, entirely new, easy-to-use, programming technique was developed specially for the COMSTATE family. This programming technique follows simple protocol logic: 1) look for conditions 2) take actions, and/or 3) go to another state. A state clearly defines all expected inputs and related actions. Protocol specific inputs and their directly related actions are presented as a trigger menu.

Powerful. The COMSTATE II will look for up to 16 triggers in any one state. 128 triggers (look for, take action, new state) can be used in the same test. One COMSTATE trigger may equal many lines of complex code in other instruments.

Engineering Applications. The power of the COMSTATE data analyzers makes them ideal for both hardware and software development. New prototypes and software can be tested without having to connect to an on-line system. New products will be developed faster and with less post-production debugging required. You will be able to: 1) test all normal operating parameters, 2) force errors to test fault tolerance and error recovery, 3) document test results, and much more.

Call immediately to schedule the most rewarding hour you can spend to increase your productivity.




```
*EM DCE/LINE*    BLOCK=225  01/17 1521
EBCDIC/8/NONE/SDLC

              TEST A                              TEST B
 * LINKUP:                          LINKUP:
 *          INIT LINK - SNRM                    LINK DEVICE UP
 * ACTPU :                          ACTPU :
 *          ACTPU TO CONTROLLER                 ACTPU COMPLETE
 * ACTLU :                          ACTLU :
 *          ACTIVATING LU                       ACTLU POS RESPONSE
 * BIND  :                          BIND  :
 *          BINDING DEVICES                     BIND POS RESPONSE
 * SDT   :                          SDT   :
 *          START DATA XFER                     SDT POS RESPONSE
 * MENU  :                          LU-LU :
 *          SEND MENU TO SCREEN                 IDLE LINK
                                  * ECHO  :
                                  *           PA2 MENU REQUEST
 * PA2   :
 *          SEND PA2 MENU           * LU-LU :
                                  *           LINK IDLE
                                  * ECHO  :
                                  *           CLEAR FUNCTION REQ
 * CLEAR :
 *          SEND ERASE WRITE        LU-LU :
                                  *           ATTENTION REQUEST
 * UNBIND:                          UNBIND:
 *          SEND UNBIND                         UNBIND POS RESPONSE
 * DACTLU:                          DACTLU:
 *          SEND DACTLU                         DACTLU POS RESPONSE
 * DACTPU:                          DACTPU:
 *          SEND DACTPU                         DACTPU POS RESPONSE
 * LINKDN:                          LINKDN:
 *          SEND DISCONNECT                     UA RESPONSE
 *          TEST COMPLETE
01/17 1523   109423    25478B       895      23       127      4257
 * LINKUP:                          LINKUP:
 *          INIT LINK - SNRM                    LINK DEVICE UP
 * ACTPU :                          ACTPU :
```

The COMSTATE II provides CRT trace and date/time stamped printouts of actual test sequences for accurate documentation of tests and results.

ATLANTIC RESEARCH CORPORATION
Teleproducts Division

7401 Boston Boulevard/Springfield, Virginia 22153
Telex: 197733 ARC TP TWX: 710-832-9828

Call Now (703) 644-9190

CIRCLE 39

Single Board Computer



6 MHz Z80

Fast, Powerful and Compact Microcomputers

- 4 or 6 MHz Z80* CPU
- 2K to 64K ROM or EPROM
- Expandable RAM from 64K to 256K
- DMA for Floppy and Hard Disk Data Transfers
- 2/4 RS-232 Serial Ports
- Centronics Parallel Port
- 50 Pin Expansion Bus for Additional I/O Capability
- CP/M Operating System and Networking Available
- Custom Systems and Private Label Packaging Available

Davidge

P.O. Box 1869
Buellton, California 93427
(805) 688-9598

*Z80 is a registered trademark of Zilog

CIRCLE 40

(continued from page 65)

detailed methodologies for testing various types of PLDs, including programmable and field-programmable logic arrays. This session will consider the application of PLD testing methods in a variety of circuits, as well as the hardware and software necessary to implement those techniques.

Another Midcon Wednesday morning session of interest to engineers concerns communication network designs. Emphasis will likely be on applications for business systems but the network design technology discussed should be universally valuable. This session is backed up in the early afternoon by one on the Integrated Services Digital Network. System component/equipment design will play a large part in the discussion but emphasis will be on the use of semiconductors to solve communication system design problems.

Late afternoon Midcon sessions rounding out the Wednesday program for designers include one on the latest in nonvolatile memory developments and another on emi/rfi. In the memory session, speakers from companies such as Intel, Monolithic Memories, and Mostek will discuss the features and applicability of EPROMs, EEPROMs, bipolar PROMs, NOVRAMs, and CMOS static RAMs, as well as battery backup where applicable. The second session will cover both the regulations involving emi/rfi and the problems involved in the compliance to those regulations. Included will be discussions on testing for FCC compliance and UL regulations, as well as the use of properly designed enclosures and shielding with conductive plastics, paint, and plating.

Thursday again will include repeats of the speech technology and surface-mounted technology Premier Sessions. In addition, there will be a potentially valuable morning session on digital signal processing and an afternoon session on computer aided engineering.

The digital signal processing session will focus on applications for presently available chips. This discussion will include comparisons of different architectures, technologies, and expected performances in specific uses. Two of the speakers will cover generic approaches in the use of these chips to implement the designs.

Recent advances in CAE technology will be the subject of the other Thursday session of specific interest. Speakers will discuss solutions to problems in design modeling, simulation, and validation—as well as the benefits offered to system designers.

All of the sessions scheduled for this year's Mini/Micro Southwest are of specific interest to computer system designers. Emphasis is on technology throughout, whether the particular subject is data communications, peripherals, or microprocessors.

Mini/Micro conference sessions

Data communications will be stressed on Tuesday, with two of the three sessions on a phase of that subject. The early morning session will be on low speed modem integration alternatives, but aimed at the design engineers who may not be aware of all facets involved. Intent of the session is to provide these engineers with the information necessary to understand the choices available to them when designing communication systems. The session will also include a review of the available LSI modem components—with emphasis here on applications.

The second communication session, in the early afternoon time slot, will examine and compare serial bus systems and protocols, and the interfacing techniques involved. Varying viewpoints will be supplied by speakers from Motorola, Hewlett-Packard, Signetics, and RCA Solid State Division.

A third communication-related session will be held in the late afternoon on Thursday. This

(continued on page 69)



NO OTHER GANG/SET PROGRAMMER GIVES YOU SO MUCH POWER FOR UNDER $2000

A powerful choice. You can use the GP256 with a dumb terminal or CPU, or as a stand-alone tool. Either way it is capable of programming most JEDEC-compatible byte-oriented EPROMs and EEPROMs on the market, from 16 to 256K bits.

Terminal Power. The GP256 converts a dumb terminal into a versatile tool for gang and set programming. Contents of multiple devices can be combined in its on-board 256K bit buffer memory. Programming time is dramatically reduced through enhanced programming methods — Intel®/Fujitsu® algorithms. SET programming permits different data sets to be simultaneously programmed into different devices. 16-bit data can be SPLIT into odd-and-even byte addresses, or conversely combined through SHUFFLE for uploading to a host CPU. Data can be edited with memory modify, insert, delete, transfer, etc.

Stand-alone power. Even as a stand-alone device, the GP256 is remarkably powerful. It can gang program up to 8 devices or set program up to 4 sets, with exceptional speed and accuracy. With a device in the socket, a few simple keystrokes issue all necessary instructions. Functions include Blank Check, Copy, Program and Verify.

Applications power. Development engineers find the GP256 useful for combining data from different files for programming, or using multiple sockets as masters when copying. Set copying speeds debugging and evaluation of system firmware. In fact, its speed and range of commands make the GP256 flexible enough to meet nearly any programming need.

And it's simple to use. So simple that operators can be trained in a matter of minutes. And multiple self-diagnostics guarantee accurate and dependable operation.

Call or write today for full details. Ask about our low-cost companion benchtop RAM and EPROM test system, the M256. And the MATE TC2000 TTL/CMOS Functional Tester or the MICE II microprocessor in-circuit emulator. They all feature the same cost-effective versatility.

Outside U.S.A. & Canada
MICROTEK INTERNATIONAL, INC.
2-1 Science Road 1
Science-Based Industrial Park
Hsinchu, Taiwan, 300, R.O.C.
Telephone: (035) 772155
Telex: 32169 MICROTEK

U.S.A. & Canada
MICROTEK LAB, INC.
17221 South Western Avenue
Gardena, CA 90247
Telephone: (213) 538-5369
Telex: 696334 BENNY GDNA

CIRCLE 41

(continued from page 67)

session will present concepts on the use of Multibus for high performance systems. Speakers will concentrate on 16/32-bit processor systems. Further emphasis will be on applications for multiprocessor systems to back up some of the other sessions in this conference on that subject.

Tuesday's final Mini/Micro technical session will be on CMOS single-chip microprocessors, with stress on performance characteristics. Further detail will be on the expected benefits of CMOS and the new application areas that are becoming available to designers, especially in industrial activities such as automotive.

Further discussion of microprocessors will be included in the Thursday morning session on cost-effective microcomputer development. The speakers promise to provide insight into alternative development tools. Both traditional and nontraditional approaches will be included in the review. The key point that will be made is that it is often possible to develop microprocessor systems without spending great sums of money.

The Thursday early afternoon session also will be microprocessor-related. This one will provide an update for system designers on available development tools. Discussions will include both software and hardware aspects.

Multiprocessor systems will receive detailed coverage in two Wednesday afternoon sessions. The early session will show how multiprocessor systems provide the additional computer power often needed for complex applications. Stress will be on realtime and fault-tolerant systems.

Increased system throughput that results from multiprocessing and coprocessing will be the theme of the second Wednesday morning session on multiprocessors. This session will emphasize applications involving both 16- and 32-bit microprocessors. Part of the discussion will cover architecture and performance of selected floating point coprocessors.

The remaining Mini/Micro technical session, this one early on Wednesday morning, will be on developments in the VLSI bit-mapped graphics controllers that are being used in so many applications. Concentration of speakers will be on graphics controllers *per se,* especially their architectures and features. Part of the discussion will, of course, include applications.—*S.F.S.*

Midcon/84 and Mini/Micro Southwest are cosponsored by Regions 4 and 5 (Chicago and Dallas Sections) of the Institute of Electrical and Electronics Engineers. For further information, contact Kent Keller, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965.

MILES AHEAD IN ENGINEERING, SERVICE AND QUALITY



Pioneer in power distribution for over 20 years. When it comes to ultra-reliable laminar bus bars for electronics, nobody comes close to Rogers for innovative engineering, comprehensive service and quality.

The perfect power distribution package. Rogers laminar bus bars are installed correctly the first time and provide better value and higher reliability than bulky wiring harnesses. (Not one of our bus bars has ever experienced an electrically-related field failure.) They fit space requirements exactly, feature low inductance, low impedance and high distributed capacitance.

Fully customized designs. Rogers engineers bus bars in a wide range of materials, tab and body configurations while offering full production runs, complete prototyping capabilities and application engineering expertise.

Worldwide service. Rogers facilities in the U.S., Japan, France and Belgium offer customers around the world the industry's best combination of innovative engineering, quality, price and delivery.

Free design/value analysis. To find out how our laminar busses, MINI/BUS® PCB bus bars and QUIK/BUS® do-it-yourself bussing can give you a head start on packaging reliability, call the Rogers Power Distribution Engineer today at (602) 830-3370.

ROGERS

Rogers Corporation
Bus Products Division
5750 East McKellips Road, Mesa, AZ 85205
602 830-3370

IT TAKES A LOT OF DRIVE

In CAD/CAM. In seismic and transaction-based systems. Or PBX's. Or anywhere else you need a lot of data fast.

That's why we created the Ampex 825.

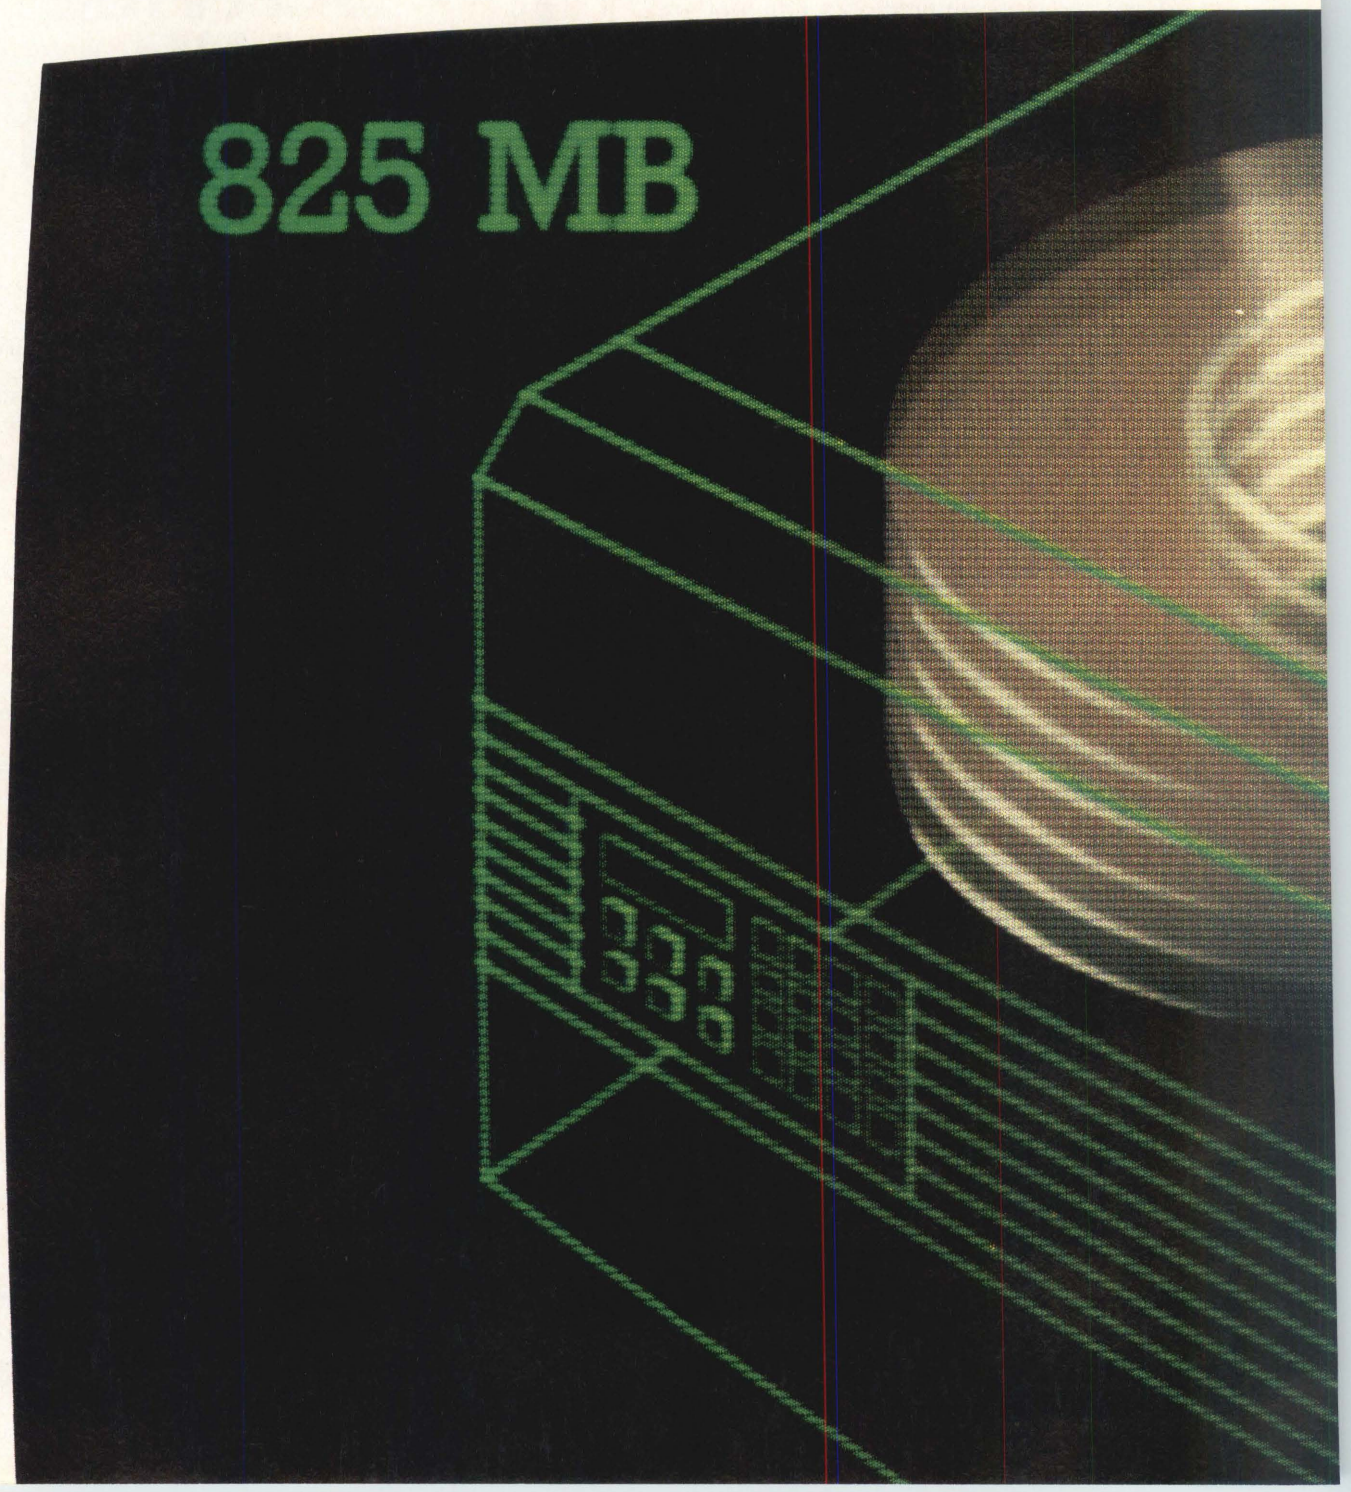
This 14" Winchester family delivers everything you need to take the lead in multi-user, multi-tasking applications: 825, 660 or 330 MB of unformatted capacity.

The new standard 1.859 megabyte per second data transfer rate.

Access times of 21 milliseconds.

Plus the best price per megabyte on the market.

Fact is, we offer a better 825 MB price than any combination of small drives. And–thanks to a design that teams RLL encoding with proven, standard head and media technologies–better yields than other big drives. So you won't have to wait for us to get our act together in manufacturing.



TO MAKE IT BIG IN PICTURES.

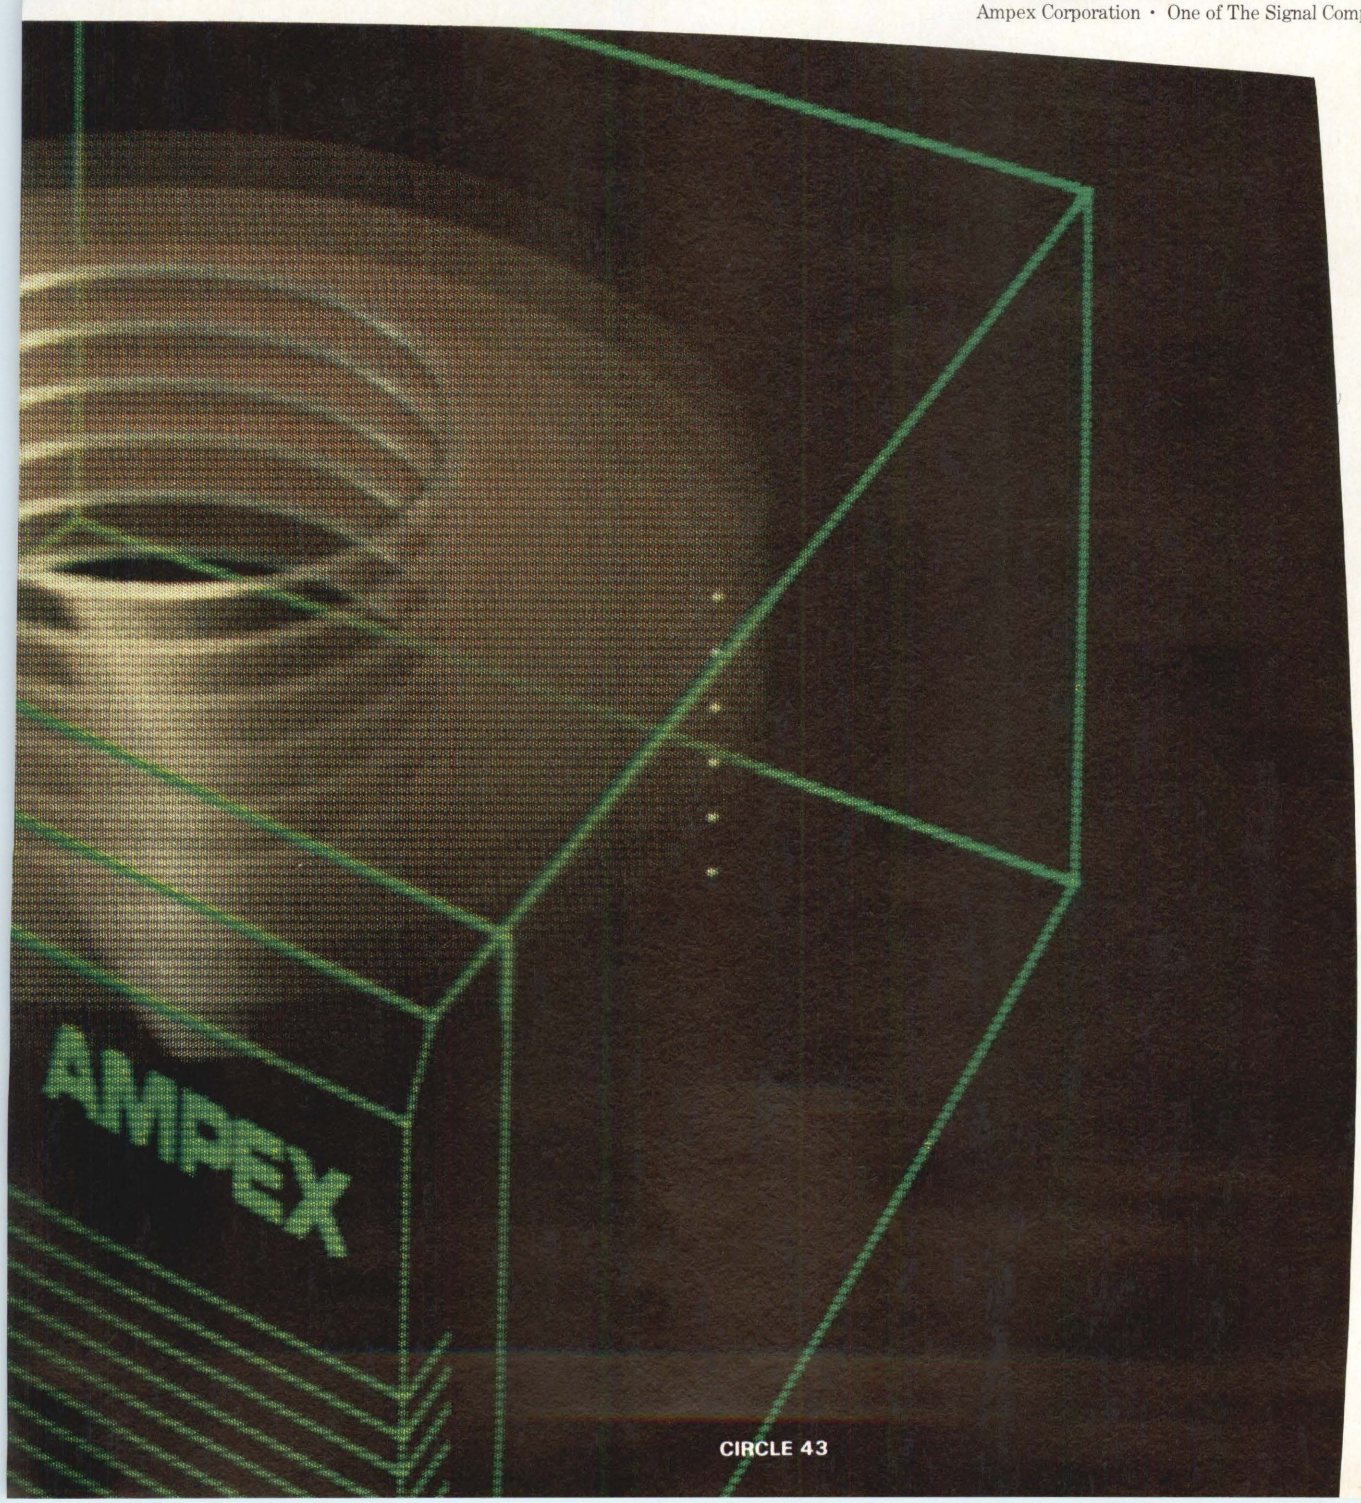
Or in QA. Or field service. Because we designed the Ampex 825 with features that keep it on the fast track. There's a unique control panel with an LED/keypad for push-button diagnostics and configuration. A modular head/disk assembly so you can upgrade capacity in minutes. A reliable, brushless DC motor, a universal power supply and a dual port option.

Plus something else no one else can match: our vertical integration and 20 years of computer peripheral and offshore manufacturing expertise.

So if you're looking for someone with the drive it takes to play a supporting role in your next big release, contact Ampex. Call us tollfree at 800 621-0292. 800 821-9473 in California. We'll be happy to set up an audition.

AMPEX

Ampex Corporation · One of The Signal Companies



CIRCLE 43

IS YOUR CHIP MAKER



Am8052
Most advanced
CRT Controller.

Am7910
First multi-spec
single chip modem.

Am29116
Fastest 16-bit
microprocessor.

Am29540
First bipolar
LSI FFT processor.

Am7990/91A
First LSI
Ethernet chip set.

Am29818
First IC
with diagnostics.

Am7901
First programmable
CODEC/Filter.

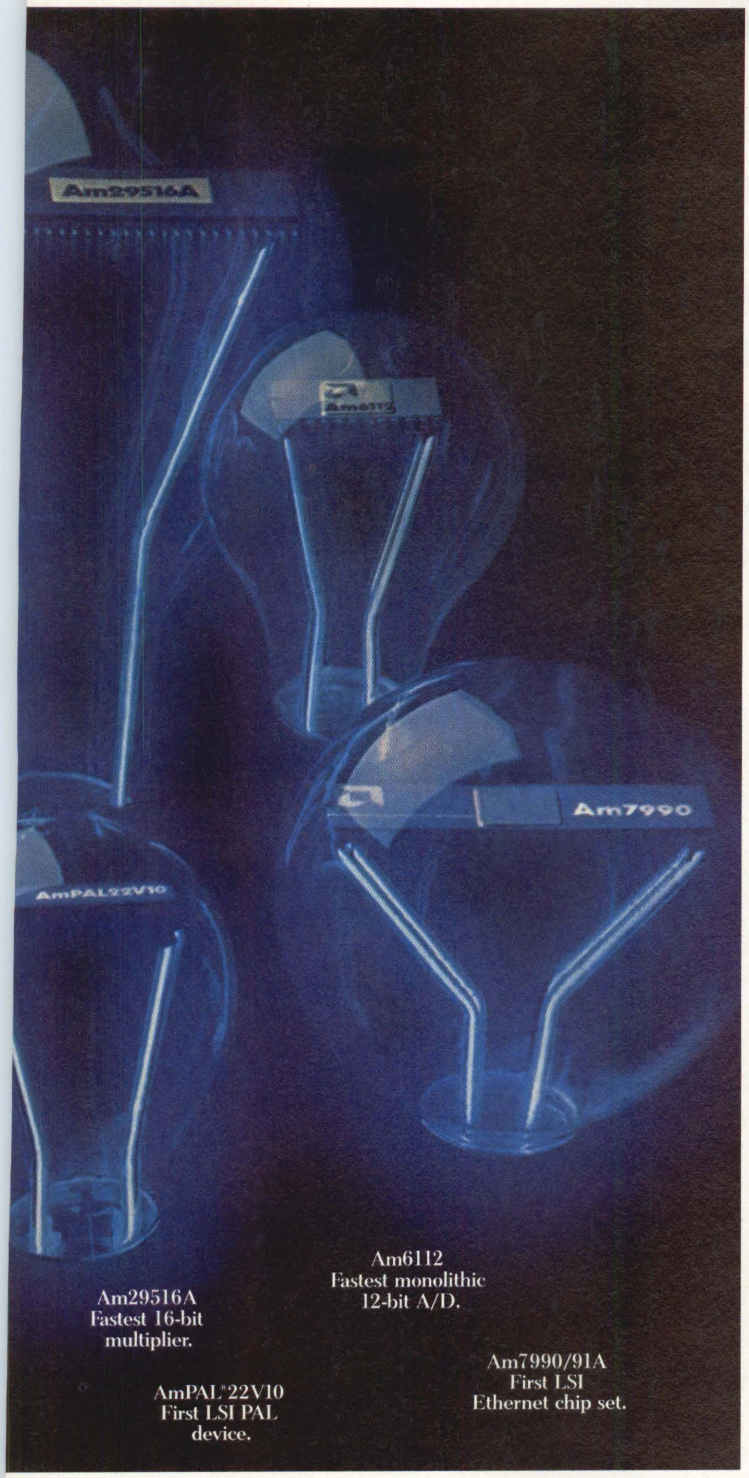
Am27S43A
Fastest 32K
bipolar PROM.

Am9150
First clearable
static RAM.

Am27512
First
512K EPROM.

PAL is a registered trademark of and is used under license from Monolithic Memories, Inc.,
© Advanced Micro Devices 1984.

AS BRIGHT AS YOU ARE?



You're doing everything you can to stay ahead of the competition. Shouldn't your IC company do the same?

We think so.

In 1983 we spent a record-breaking 18.7% of sales on research and development. That's more than any other major IC company.

And that's why 40% of our total sales come from products that were invented here.

Our Am8052 CRT controller lets you squeeze every last drop of performance out of a video tube.

And our Am7910 is the only modem you can program for any major telephone system anywhere in the world.

We make the first perfectly matched VLSI Ethernet chip set and the only complete kit solution to super high speed digital signal processing.

We make the world's first 512K EPROM and the world's fastest microprocessor.

We're even an innovator when it comes to quality.

While other guarantees run on and on, ours is short and sweet.

The International Standard of Quality guarantees a 0.1% AQL on all electrical parameters, AC and DC, over the entire operating range.

INT-STD-1000

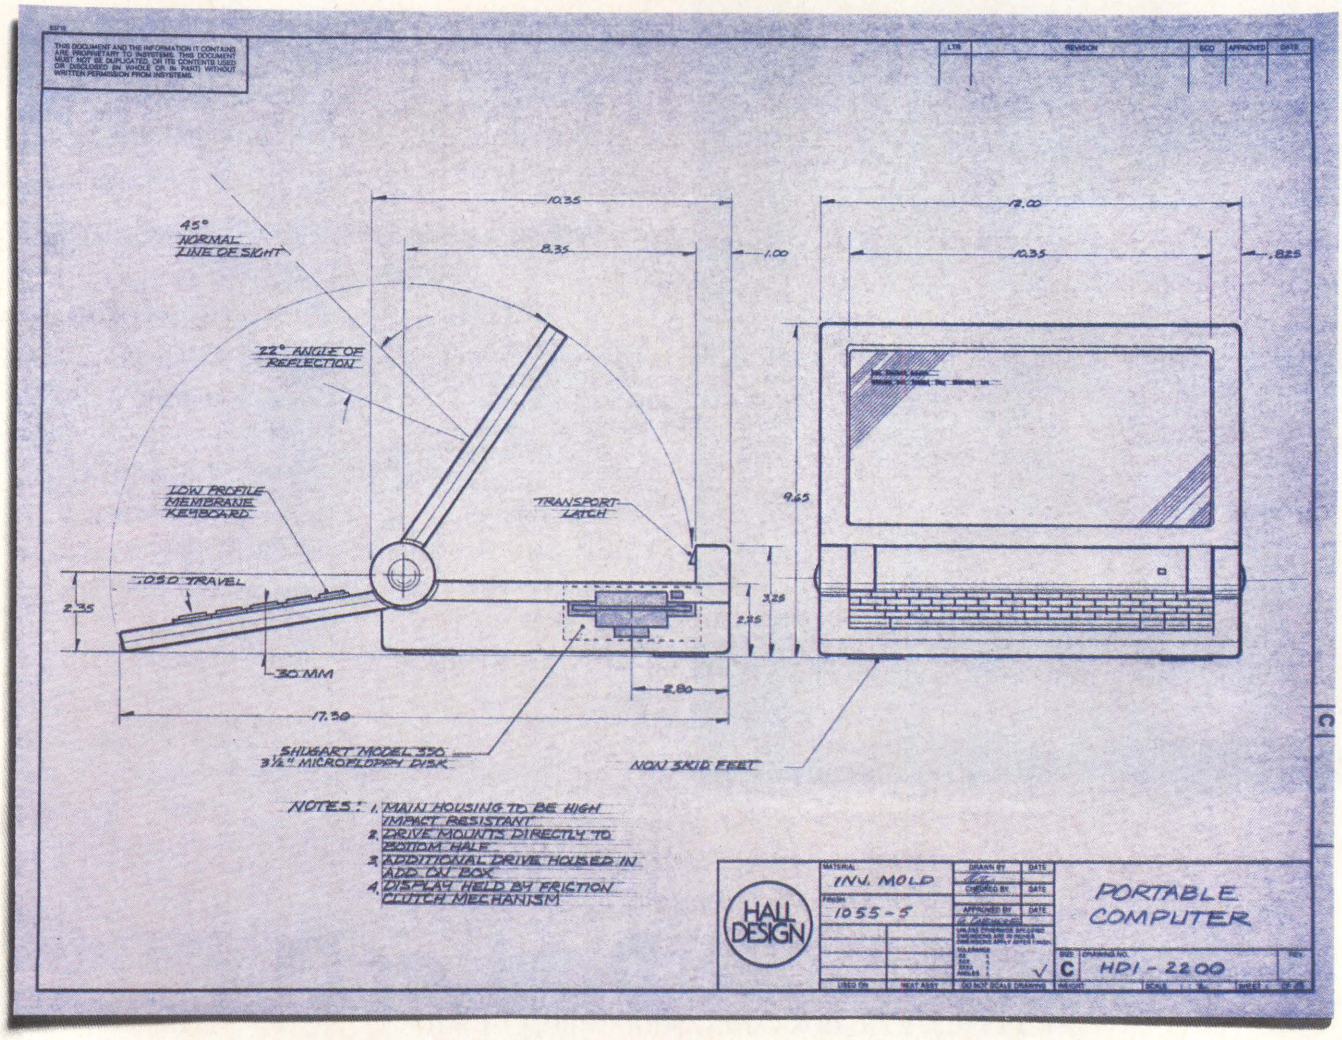
So, if you've got a product you want to really shine, call us. We can put you as far ahead of the competition as we are.

Advanced Micro Devices

901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088
(408) 749-5000, outside California, call toll free (800) 538-8450, ext. 5000.

CIRCLE 45

IF YOUR BUSINESS IS SHRINKING, SHUGART CAN HELP.



These days, computer designers everywhere face a problem of massive proportions: How do you cram a desktop computer into a briefcase?

Sound familiar? If so, there's a family of 3.5" single and double-sided microfloppy disk drives you should meet. The Shugart 300 and 350, respectively.

The perfect drive solution for a full-featured portable.

Tiny enough to fit easily into your smallest design. Yet with a 6 millisecond average access time, a capacity up to 1 megabyte and Minifloppy™ compatibility, your portable computer could easily run the same software as someone else's desktop.

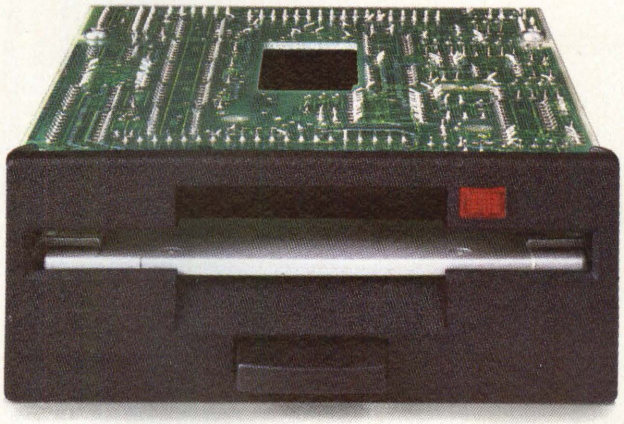
And keep it running for quite some time.

Shugart drives are so reliable, you can count on an MTBF of 10,000 power on hours. One reason we're projecting delivery of over 100,000 microfloppy drives this year.

At just over a pound apiece, you could even use two. And still call your portable computer portable.

There's just one thing to remember. You should always check the activity light on a Shugart microfloppy. They're so quiet, there's no other way to tell if they're running.

What more could a shrinking business need? A couple of other small things. Industry standard 3.5" microcartridges, to be exact. Their track densities offer a more than generous upgrade path. But more important, considering where they could end up, they come equipped with a hard shell plastic media cartridge. And an automatic head access shutter. Sure protection from all kinds of catastrophes. Stick them in your pocket. Throw them in your purse. Bang them around in your briefcase. They'll survive.



The 3.5" Shugart Microfloppy. Smaller than actual size, but not much.

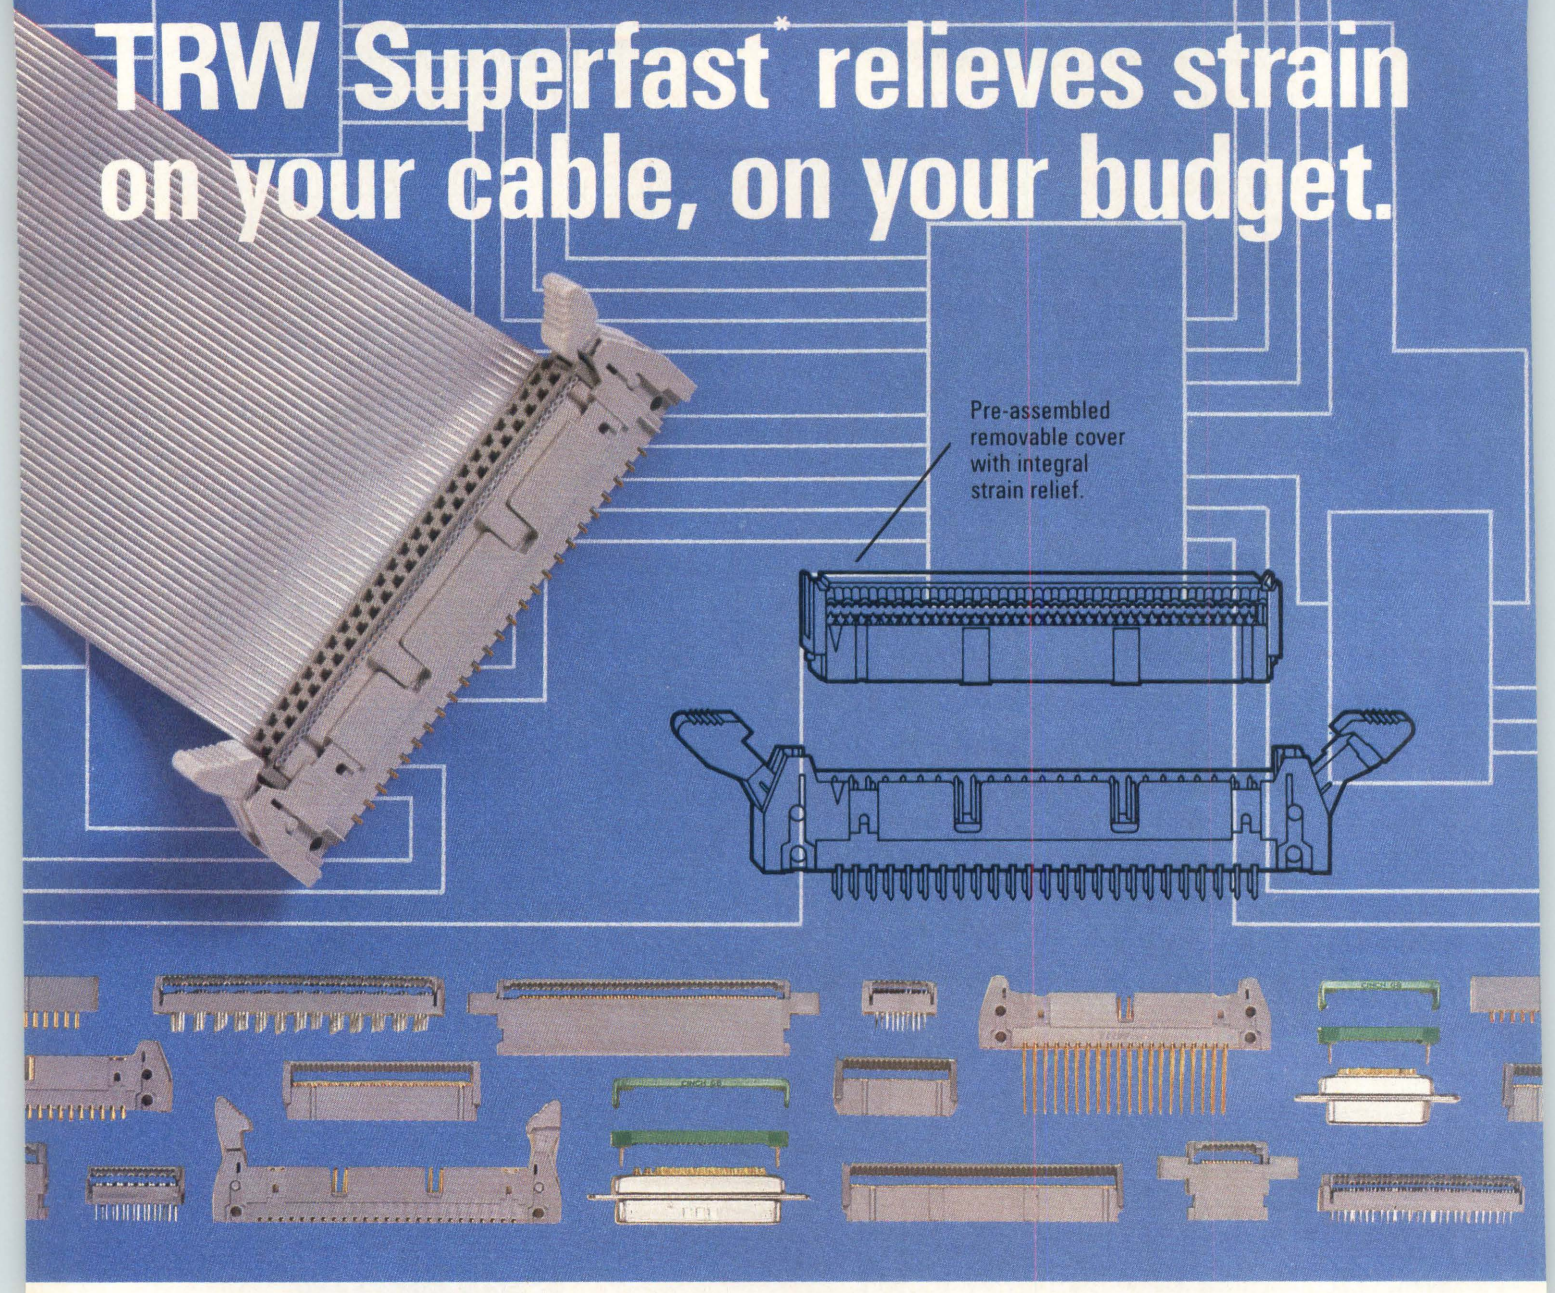
Sound interesting? Call your local Shugart sales office. We'll do a private microfloppy workshop right in your office. But do it soon. The more your business shrinks, the more Shugart can help.

Shugart

Right from the start.

Milpitas, CA (408) 263-2600; Costa Mesa, CA (714) 979-1935; Thousand Oaks, CA (805) 496-5388; Rochester, NY (716) 235-7190; Minneapolis, MN (612) 546-4411; Richardson, TX (214) 234-3568; Framingham, MA (617) 879-1700; Saddle Brook, NJ (201) 368-8445; Smyrna, GA (404) 436-0953; Markham, ONT (416) 475-2655; Paris, France (3) 946-42-66; Munich, West Germany (089) 786-021; London, U.K. (44) 4862-27272; Wanchai, Hong Kong (852) 5-733307. © 1984 Shugart Corporation

TRW Superfast* relieves strain on your cable, on your budget.



Specify TRW Superfast Flat Cable Connectors

This is the flat cable mass termination interconnect system that's designed to relieve strain in every way.

The cover/strain relief is pre-assembled as an intergral part of the connector. So you're assured of effective cable strain relief without extra parts to buy or assemble. And one, simple, snap-action operation puts it in position.

TRW Superfast connectors are just as easy on your budget. Low prices to begin with. More density because of the lower profile of the overall package height. And the Superfast design requires less cable, less labor and reduces space requirements when used in a daisy chain mode. Superfast even meets MIL-SPEC requirements.

All this is part of the Value Engineering that TRW built into every Superfast Flat Cable Connector. It also includes features like an exclusive one-step insulation displacement contact system with a unique contact design for high reliability connections.

And you can get all the connector sizes and types you need in the Superfast system. More important, you can get fast, off-the-shelf delivery from the TRW Connector distributor network—the best in the business—to prevent strain on *your* commitments.

For more information on the Superfast Flat Cable Connector System, contact your TRW Electronic Components Group sales office, or see your local TRW Connector Distributor. Or write: Connector Division, TRW Electronic Components Group, 1501 Morse Avenue, Elk Grove Village, Illinois 60007. Phone 312.981.6000.

*Trademark TRW Inc.

©TRW Inc. 1984.

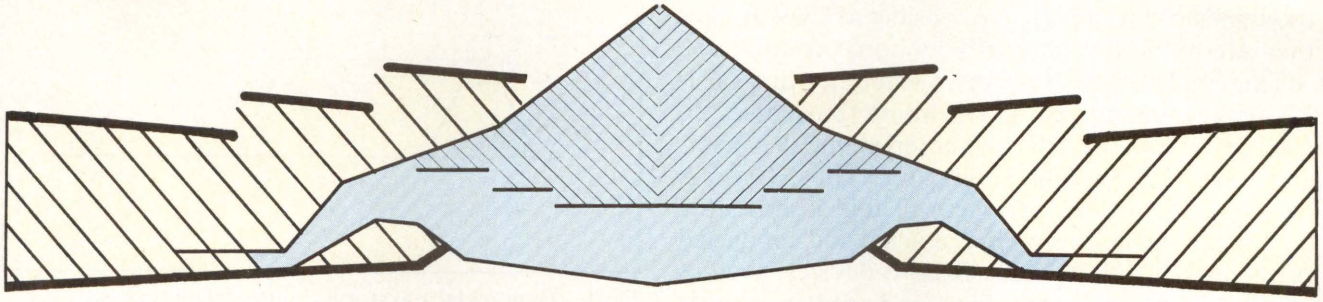
CIRCLE 47

TRW

Connector Division
TRW Electronic Components Group

CONNECTORS—THE MISSING LINK IN EMI SUPPRESSION

Computer equipment must meet stringent FCC restrictions on rf emissions. Although the first line of defense for designers is proper packaging, electromagnetic interference can occur between units if attention is not paid to connector design.



by Frank Drzymkowski and Dave Goodman

Since the early days of radio, communication engineers have been concerned with electromagnetic interference. Until recently, however, digital system designers did not have to pay much attention to electromagnetic interference. This was because computer systems were large and generally isolated in their own rooms. Even if they did radiate electromagnetic interference, there was nothing else close enough to be affected. Now, with the proliferation of computing systems and their penetration into small business and home environments, system designers can no longer afford to ignore this aspect of computer design.

Recognizing the potential interference problems represented by the rapid spread of digital systems, the Federal Communications Commission has imposed strict regulations governing the amount of electromagnetic interference (emi) that a digital system may emit. The FCC docket #20780 covers computers and all other electronic devices that generate and use radio frequency (rf) energy for timing and control applications. For the purposes of this regulation, rf energy includes any signal with a frequency between 10 kHz and 1000 MHz. As of Oct 1, 1983, any equipment under the jurisdiction of this docket that does not meet the requirements cannot be sold or offered for sale. If already sold, the equipment must cease operation until brought into compliance. Fines of up to $2000 a day can be levied for noncompliance.

The new regulations are concerned with only two of the several forms of emi: radiated and conducted emissions. Other types of emi, such as ac hum and electrostatic discharge, are not considered. The FCC specifies the test methods that must be used to verify compliance, as well as the permissible levels of emitted emi.

Frank Drzymkowski is director of engineering at ITT Cannon Phoenix, 2801 Air Lane, Phoenix, AZ 85034, where he is responsible for product design and manufacturing/engineering. He holds a BEE from the University of Dayton and an MSEE from Purdue University.

Dave Goodman is director of new product development engineering for ITT Cannon North America, 10550 Talbert Ave, Fountain Valley, CA 92708. Mr Goodman holds a BSME from the Georgia Institute of Technology.

TABLE 1
Class A Equipment emi/rfi Limits

	Frequency (MHz)	Distance (m)	Field Strength (V/m)
Radiated:	30 to 88	30	30
	88 to 216	30	50
	216 to 1000	30	70
	Frequency (MHz)	Maximum Voltage (V)	
Conducted: (power line)	0.45 to 1.6	1000	
	1.6 to 30	3000	

Docket #20780 divides equipment into two classes: Class A includes all devices marketed primarily for use in a commercial environment; Class B covers all devices marketed for home use. Class A devices include business computers, numerical control machines, and quality control equipment. Computers such as Apple, even though they are used in thousands of businesses, are considered Class B since they are marketed primarily as home computers.

There is less potential for interference in a business environment than in the home (businesses do not have the wide variety of entertainment systems and other devices that homes do). Thus, Class A requirements are less stringent. Table 1 shows the FCC requirements concerning radiated and conducted emi from Class A devices. Class B devices, such as home computers, videogames, and cordless phones are used in smaller quarters as well as in close proximity to televisions and radios. Therefore, requirements are more stringent. Table 2 shows the FCC limitations on radiated and conducted emi from Class B devices. Note that while Class A limits are specified at 30 m, Class B limits are specified at 3 m. This reflects the closer spacing between potentially interfering sources in the home.

Perhaps the most important aspect of the FCC regulation is that it addresses system design rather than the components or the technology chosen to build a device. While the docket states that the manufacturer is responsible for meeting the radiated and conducted noise restriction requirements of the regulation, it does not require any specific design or component to be used. This leaves the system designer free to choose the most efficient and economical method of controlling emissions for a particular application.

TABLE 2
Class B Equipment emi/rfi Limits

	Frequency (MHz)	Distance (m)	Field Strength (V/m)
Radiated:	30 to 88	3	100
	88 to 216	3	150
	216 to 1000	3	200
Conducted: (power line)	At any frequency from 450 kHz to 30 MHz, conducted rfi and emi must not exceed 250 V.		

Digital systems and emi

In any electronic system, radiated noise is likely to be the dominant form of interference. Virtually any part of a circuit, whether a trace on a PC board or a length of discrete wiring, can act as a broadcast antenna. Conducted emission of rf noise usually becomes a problem when it reaches a part of the circuit where it can be radiated. These problems become more pronounced as the frequency increases. Conducted emi can also be a problem in equipment with switching power supplies, where leakage back into the power lines can cause interference in any piece of equipment not protected by adequate filtering.

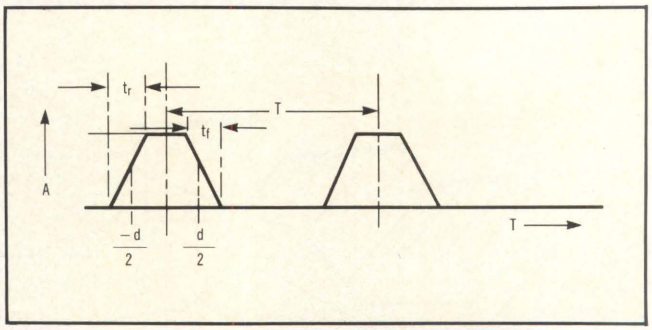


Fig 1 Trapezoidal waveforms approximate signals. Controlling emi is difficult due to the large resource of harmonics found in digital control signals because of fast rise times and high frequencies.

Unfortunately for the system designer, digital electronics are noisy by nature. The fast rise times and high frequencies characteristic of digital control signals make them rich sources of harmonics, which complicates emi control. The trapezoidal waveform (Fig 1) is a good approximation of the signal type found in computing devices. In the Fourier analysis of this waveform, as shown below, A is pulse amplitude in volts, d is pulse width in seconds, n is a harmonic number, T is pulse period in seconds, and t_r is rise time in seconds.

$$e(t) = \frac{Ad}{T} + \frac{2Ad}{T} \sum_{n=1}^{\infty} \frac{\text{Sin } n\pi t_r/T}{n\pi t_r/T} \frac{\text{Sin } n\pi d/T}{n\pi d/T}$$

$$n = 1, 2, 3, \ldots$$

This analysis yields the spectrum envelope shown in Fig 2 (the rise and fall times of the signal are assumed to be equal).

The control signals used in the three major families of logic—CMOS, TTI, and ECL—are each characterized by different voltage swings and rise times. Using a 5-MHz version of the waveform in Fig 1 as the basis for Fourier analysis of each of the logic types, the spectrum envelopes illustrated

in Fig 3 are obtained. Note that all three types of logic generate appreciable harmonics out to the gigahertz region. At the higher frequencies, where shield integrity and conductor coupling become more important, TTL generates the highest levels of harmonics. This worst case will be used as the model for discussion.

An antenna model, with a 20-cm length of 20 AWG wire 5 cm above a ground plane, and carrying a current of 10 mA, will be used to estimate how much radiated interference a digital circuit using this signal might generate. This approximates some of the longer wire traces on a single-board computer, for instance. Fig 4 shows the result of using the standard Maxwell field equations, and taking into account the directional gain of such an antenna. The signal radiated by the model is at least 10 dB higher than the FCC allows for Class B digital systems. [A detailed treatment of the antenna theory behind this graph is available from the authors.] Although the exact circumstances represented by this model might never be encountered in digital system design, it nevertheless illustrates the kind of problem a digital engineer should recognize.

A designer can choose between several methods to deal with emi: case shielding, onboard filtering, multilayer PC boards, etc. Case shielding is a *sine qua non*—no attempt at emi control will get very far without it. Multilayer PC boards can be used to furnish a ground plane very close to any possible antenna element, thus limiting significant radiation to the very highest frequencies. This approach has two disadvantages: it is somewhat expensive; and when system redesign is necessary to overcome emi problems, it is rarely appropriate. Conducted radiation can often be dealt with most simply by providing a filter network at the output of the board or subsystem generating the offending signal.

Dealing with emi

A good design uses one or all of these methods as the application dictates; none are sufficient by themselves. A case may provide a perfectly seamless shield, eliminating all radiated energy, but it will not affect conducted noise. Thus, any rf noise generated in a perfectly shielded computer connected to an imperfectly shielded peripheral (and this includes the power lines), will be conducted to the peripheral and be radiated through the imperfect shield. A related problem is the expense of well-shielded cable that may be needed to connect subsystems in a digital design.

Filtering the signal to eliminate unwanted rf components eliminates conducted noise, but without a well-shielded design, radiation still occurs. In some cases, when the signal itself is the source of noise rather than its harmonics, the signal waveform will determine how easy it is to filter. A

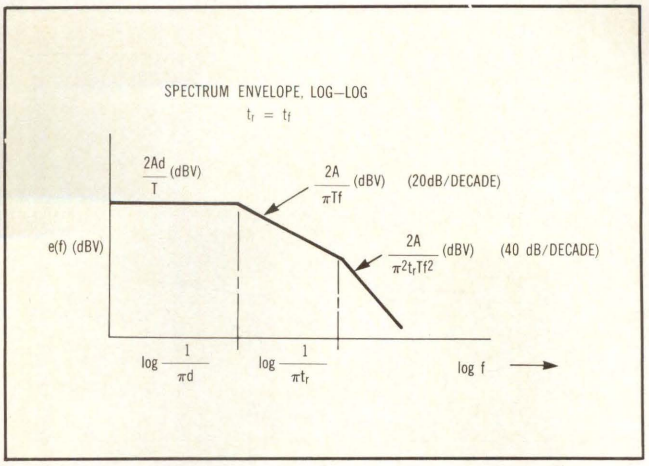


Fig 2 The spectrum envelope shown is developed from a Fourier analysis of the waveform in Fig 1. Here, the signal's rise and fall times are assumed to be equal.

sinusoidal waveform is more easily dealt with than more complex signals.

Signal frequency and power strongly affect the efficacy and cost of each approach. The higher the noise frequency, the smaller the aperture through which it can escape, and the harder it is to shield. Thus, digital systems require a case that is as seamless as possible. The power level of the interfering signal determines how far radiated noise will travel, as well as how far it will be conducted. Filtering a high power signal is naturally more expensive.

Even a system designed with careful attention to all of these factors can be compromised by weaknesses in other system components. The design elements most often overlooked are the connectors used to link the various parts of the system. Connectors, like power supplies, cases, and other "secondary" elements of a system, are often specified only after

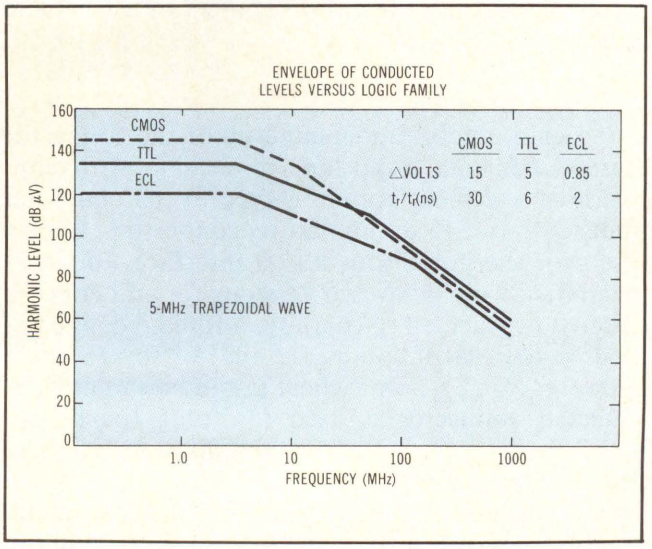


Fig 3 Fourier analysis of major logic families—CMOS, TTL, and ECL—yields the spectrum envelopes illustrated. All three logic types generate significant harmonics to the gigahertz region.

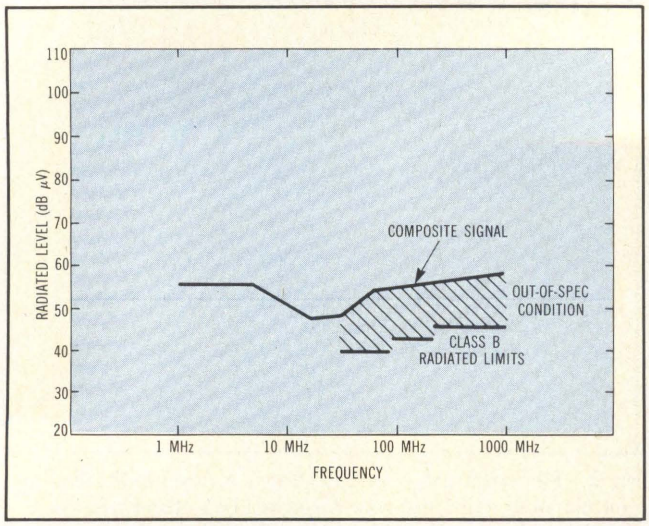


Fig 4 Maxwell field equations applied to a digital circuit antenna model yield a composite signal in excess of the FCC limits.

system design is complete—as an afterthought. As a result, much of the digital equipment in the field today does not meet FCC standards. (Spot checks by the FCC have revealed that up to 30 percent of Class B systems are not in compliance.)

Connectors specified for a digital system can play a very important part in emi control. An improperly designed connector can compromise an otherwise excellent design by allowing the radiation or conductance of internally generated radio frequency interference (rfi) and emi. If a piece of equipment must be redesigned to meet the new FCC regulations, filtered or shielded connectors may offer the most economic fix. A knowledge of the basic elements of connector design for emi control can help a digital system designer choose the connectors that will enable development of a quiet and economical system.

Shielded connectors

To gain a good understanding of the connector design elements important in dealing with emi, consider the most common type in the computer industry—the D-subminiature connector. Widely used for the ubiquitous RS-232 interface, this connector is available in two versions for emi control: filtered or shielded. Generally, a filtered connector is used to control conducted emi, or to prevent it from reaching a point where it can be radiated. A shielded connector is used to "plug the holes" in a shielded case. Both types allow a design to be brought into compliance with the FCC docket without changing the basic design. Thus, the manufacturer does not incur a great expense. In addition, the effect of new connectors on emi control is easier to assess and test than any other design change.

Ideally, when shielded connectors are mated, they should provide the same level of shielding as the cable shield and case. If not, the connector becomes a hole through which emi can escape from an otherwise well-shielded system. This requires special attention to the integrity of the connector housings, the backshell, and the joint between any two pieces of a mated connector, including the cable and the case. The method used to terminate shielded cable is especially important, due to the fact that exposed, unshielded wires at the connector termination are free to radiate. At the same time, the method chosen to ensure shielding integrity must do so without compromising ease of assembly. Another consideration often overlooked is the susceptibility of an otherwise well-shielded design to damage from electrostatic discharges.

To judge the relative effectiveness of each design element of a shielded connector, some measure of performance must be devised. It should be fairly simple and as independent of extraneous factors as possible. Cannon, for example, has chosen to rate connector designs using transfer impedance, a measure of performance that is also widely used in the wire and cable industry. Shielding effectiveness, which is an alternate criterion and a more direct measurement of performance in some ways, is not used because of its dependence on factors such as terminal impedance and the configuration of the measuring apparatus—factors that do not specifically affect shielded connector design.

Transfer impedance is the open circuit voltage induced in the internal conductors of a shield by longitudinal current flowing on the shield's exterior surface. The transfer impedance of a shielded connector can be represented by $Z_t = R_0 + j\omega M_{12}$. Here, $j = -1^2$, M_{12} is mutual inductance between two lines expressed in henrys, omega is angular frequency in radians, (equal to $2\ \pi f$), R_0 is a real component of transfer impedance in ohms, and Z_t is transfer impedance in ohms. By using transfer impedance, the relative effectiveness of each shielded joint in a connector can be judged through direct measurement, and an overall measurement by which to judge various design approaches can be determined.

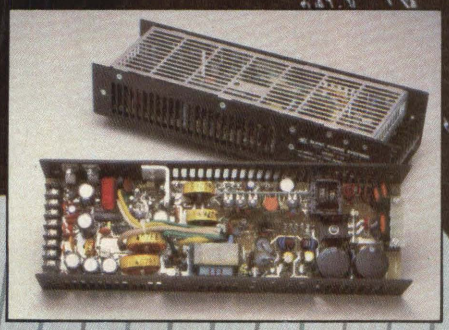
To illustrate the various aspects of shielded connector design important to a system designer, consider Cannon's shielded D-subminiature connectors. Each of the connector's parts addresses a specific aspect of emi control, with the overall goal of shield integrity. The shielded backshell of this connector is a one-piece, drawn brass design that encompasses all wiring between the cable shield and the connector shell. Fig 5 (a) shows the transfer impedance of this design versus a two-piece die-cast design. The leakage permitted by the two-piece design is approximately an order of magnitude greater than that of the one-piece backshell. This is partly due to the difficulty of producing a two-piece design whose mating edges are perfectly smooth, and

WE'LL GO HEAD TO HEAD WITH ANY 200 WATT SUPPLY!

ELPAC'S

MASS 210

THE ELPAC MAGNETIC AMPLIFIER SWITCHER SERIES



Powers Disk, Logic and CRT with the same 12v Output!

- **Provides instantaneous peak load regulation**
- **Open frame construction with protective screen reduces EMI/RFI**
- **80% efficiency**
- **Simplified secondary design for enhanced reliability**
- **Overcurrent protection on all outputs**
- **Low profile package (2.5 X 4.25 X 13.00)**

Elpac challenges you to test the MASS 210 with any 200 watt power supply you've been using. We're sure that you will agree that the MASS 210 allows FREEDOM OF CHOICE for the system designer in the fast growing hard disk plus back-up based market. One of the reasons that this is the designers choice is the unique feature of providing instantaneous peak loading regulation. Call today to inquire about your MASS 210 power supply for a 30 day free trial.

When you need Unregulated, Linear, Linear Wall-mounts, Switch Regulated, or DC/DC Converters, Elpac has a power supply that will fit your design.

	OUTPUTS			
PART NO.	**1**	**2**	**3**	**4**
MASS 210-1*	5V/15A (20A)**	12V/6A (8A pk)	−12V/1A (1.5A)**	—
MASS 210-2*	5V/15A (20A)**	12V/6A (8A pk)	−12V/1A (1.5A)**	24V/4A (6A pk)
MASS 210-3*	5V/15A (20A)**	12V/6A (8A pk)	−12V/1A (1.5A)**	−5V/5A (.75A)**
MASS 210-4*	5V/15A (20A)**	12V/6A (8A pk)	−12V/1A (1.5A)**	+28V/3A (5A pk)
MASS 210-5*	5V/15A (20A)**	12V/6A (8A pk)	−12V/1A (1.5A)**	+12V/4A (6A pk)

*TOTAL POWER AVAILABLE 210 WATTS WITH 50 CFM AIR FLOW
**REPRESENTS MAXIMUM AVAILABLE CURRENT

ELPAC - SERVING APPLICATIONS WITH A FREEDOM OF CHOICE!

ELPAC POWER SYSTEMS

3131 S. Standard Ave. • Santa Ana, CA 92705 • (714) 979-4440 • TWX 910-595-1513

ANOTHER UGLY® Where the beauty is service & performance.

GET UGLY ELPAC

©Copyright 1984 ELPAC - all rights reserved

CIRCLE 48

partly due to the formation of nonconductive (on aluminum) or semiconductive (on brass) oxidation layers on the joint edges. This latter effect plays a major role in all joints in a shielded connector, and mandates the use of some form of protective finish for best emi protection.

The cable shield is terminated to the shielded backshell by means of a thin-wall cylindrical ferrule using hexagonal crimp pliers. This shield is captured in an annular band between the outer diameter of the ferrule and the inner diameter of the shielded backshell. This technique can be used with either braided or foil shield cable. Fig 5 (b) shows the transfer impedance of this joint, including 0.5 in. of cable shield. It is interesting to note that although a foil shield with drain wire gives better shielding coverage when cable construction alone is considered, the transfer impedance of its joint to the shielded backshell is approximately an order of magnitude higher than that with braided shield cable. This is due principally to the inductance of the drain wire that furnishes the electrical connection between the cable shield and the backshell crimp.

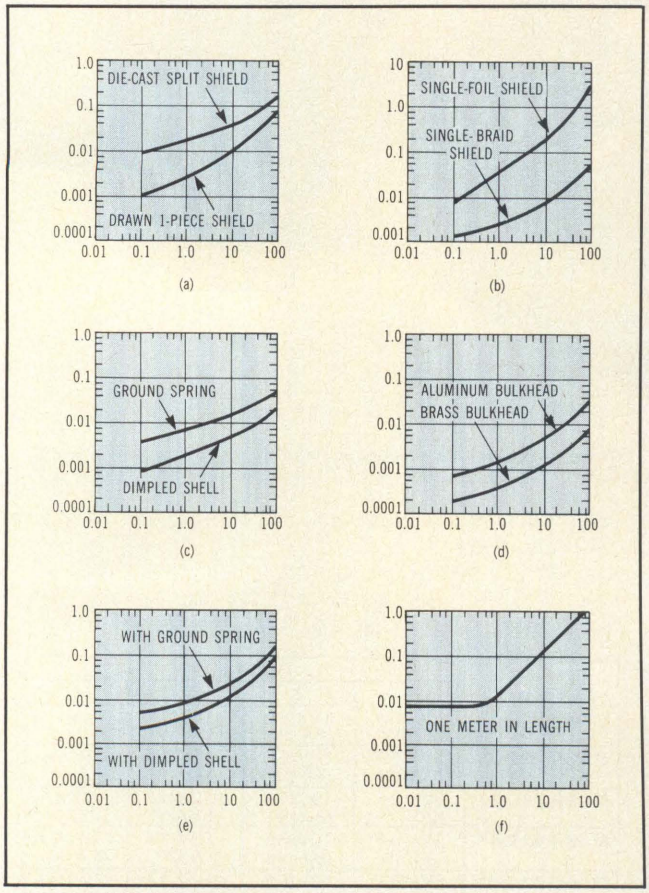


Fig 5 Pictured are the transfer impedances (ohms/ y axis) versus frequency (MHz/ x axis) between, for example, shielded backshell and connector backshell (a); between shielded backshell and cable shield, including 0.5 in. of cable shield (b); between plug and receptacle shells (c); and between receptacle and bulkhead (d). Also pictured are the total of all contributions from aluminum bulkhead through to shield crimp (e); and the typical braid shielded cable (f), which is 0.3 in. in diameter.

The steel connector shells continue the shielding through the connector receptacle mounting panel. This is an extremely important joint for controlling emi: even if the connector and socket shells overlap entirely, there is at best only sporadic electrical contact between them unless one of two methods is used to ensure a low resistance joint. One method—multiple grounding dimples on the male connector shell—furnishes a number of low resistance pathways upon engagement. One advantage of this design is that it will mate with any standard female D-subminiature connector shell. As long as the female connector is grounded to the case, adding a dimpled male connector will give a well-shielded connection—a major advantage in case of FCC-mandated redesigns.

Another method is the use of cantilever grounding fingers on socket connector halves. This approach leaves both mating shells unchanged but requires that the insulator of the socket half of the connector pair be slotted to accept this grounding finger addition. One advantage of this design is that it permits direct grounding to the traces on a PC board. Fig 5 (c) illustrates the differences in transfer impedance between these two designs. The dimpled shell enjoys approximately a 6-dB advantage, a minor difference that is overshadowed by other factors in connector design.

The final joint to be considered is that between the socket half of the connector pair and the bulkhead to which it is grounded. Fig 5 (d) shows the difference in transfer impedance exhibited by brass and aluminum bulkheads. Here, too, the formation of high resistance oxidation layers plays a part, and no amount of bearing pressure exerted between socket and bulkhead appears to have any effect. Again, protective finishes are a necessity, especially in products that may be exposed to elevated temperature or humidity.

Fig 5 (e) sums up the contributions to the connector's total transfer impedance of all these design elements, omitting only the cable shield-to-backshell junction, which is highly dependent on the cable used by the design. In comparison, Fig 5 (f) illustrates the theoretical transfer impedance/meter length of a 0.3-in. diameter cable with a shield consisting of 7 strands of 36 AWG wire in each of 24 carriers and a 30-degree weave angle. It is substantially higher than that of the entire connector.

One last part of connector design important to digital designers is the connector shroud. This insulates the shielded backshell from electrical contact with any other part of the system, or with operating personnel.

There are other design elements that should be considered in selecting a connector. The method of strain relief should not rely on the mechanical



YOU'LL NEED A PARTNER WHO CAN CUT IT.

How We Look At The Future. Designing information systems for the business office of the future is a lot like planning the flawless performance in ice skating.

Choosing the right printer partner can be critical.

Are the same strong goals for success shared? Is the necessary talent, commitment, and dedication to meeting and exceeding those goals present?

As a major designer and manufacturer of state-of-the-art printers, worldwide, Okidata knows the importance of goals and commitment. And living up to them.

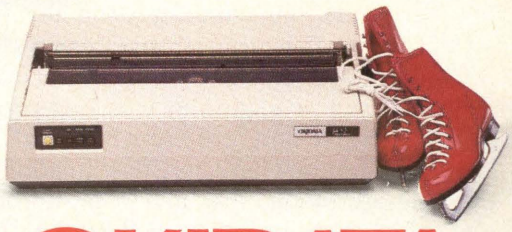
What We're Doing Today. For Tomorrow. Right now, our dedicated research and new product design teams are pushing and testing the limits of present technology to find better ways to build better printers.

Through an on-going and expensive commitment to robotic assembly, we're assuring smoother and faster-than-ever product flow.

And, elsewhere, we're streamlining our customization and modification turnaround times to respond even more quickly to your rapid startups.

We'll Be There When You Need Us. In OEM system building, just like in the Olympics, commitment is everything.

If your audience will be looking to you for more flawless performances in the future, we're the printer company who'd like to join you. In fact, we're already working on it. Call 1-800-OKIDATA. Or write OKIDATA, Mt. Laurel, NJ 08054.



OKIDATA
an OKI AMERICA company

CIRCLE 49

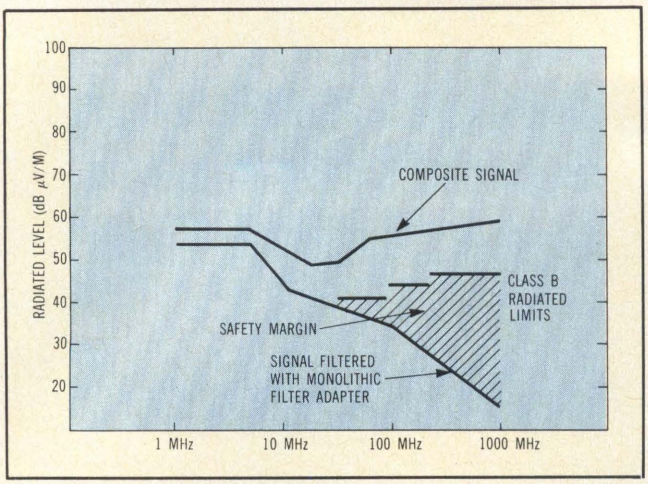


Fig 6 The addition of a filtered connector to the antenna model analyzed in Fig 4 brings radiation well within FCC specifications.

strength of the crimp, especially with foil shielded cable. It should however, accommodate a range of cable sizes. The latching mechanism that holds the connectors together should be simple and reliable, or the familiar captive screw fasteners should be used where frequent disconnection and reconnection are not anticipated. A designer may also wish to investigate the possibility that only a part of the connector needs to be replaced to bring a design up to spec. The Cannon shielded D-subminiature parts can be mated with a number of other D-subminiature products. This capability can further reduce redesign cost.

Filtered connectors

In some digital systems, the emi problem may not result from a poorly shielded case or leaky connectors, but rather from emi that is being conducted out of a well-shielded environment to where it can be radiated, or from straight conducted emi in excess of FCC-mandated levels. In either case, a filtered connector may be the proper solution.

An example of filtered connector design is the transverse monolith filtered connector. Its electrical design includes a monolithic capacitor that acts as an rf bypass, shunting all signals above a frequency determined by the combination of source and load impedances in conjunction with the capacitance to ground. This eliminates the harmonics responsible for emi. The monolithic design has several advantages over earlier designs that used discrete components. For example, it is one-piece, and thus less expensive and more reliable. It also produces better attenuation characteristics. In addition, the one-piece design means that any combination of contacts in the D-subminiature shell can be filtered by one connector design. It is smaller in length than discrete tubular components, which means it can be directly substituted for any other standard D-subminiature connector, and can be combined with the shielded connector adapter kits for added protection if necessary. The monolithic design is also less subject to shock, vibration, and other environmental effects.

Fig 6 illustrates the effect the addition of this filtered connector has on the signal from the antenna model analyzed in Fig 4. Here, the antenna model is assumed to represent a wire outside a shielded system that might allow conducted emi to radiate. With the connector in place, the emi radiated from the antenna model is now well within FCC specifications.

One last advantage of both shielded and filtered connectors to keep in mind is the ease with which they permit testing of the effect that the additional connectors may have on an out-of-spec system. A monolithic filtered adapter placed between the two halves of a connector pair reveals whether the emi problem is due to straight radiated emi (in which case the adapter will have no effect), or reradiated conducted emi (in which case the addition of the adapter will eliminate or cut down on the emi measured). This is one of the first tests that should be performed on an out-of-spec system. The same type of test can be made with a shielded connector to see if a leaky connector is permitting radiation of emi. Both of these tests can prevent needless work and help bring a system back into compliance quickly and economically.

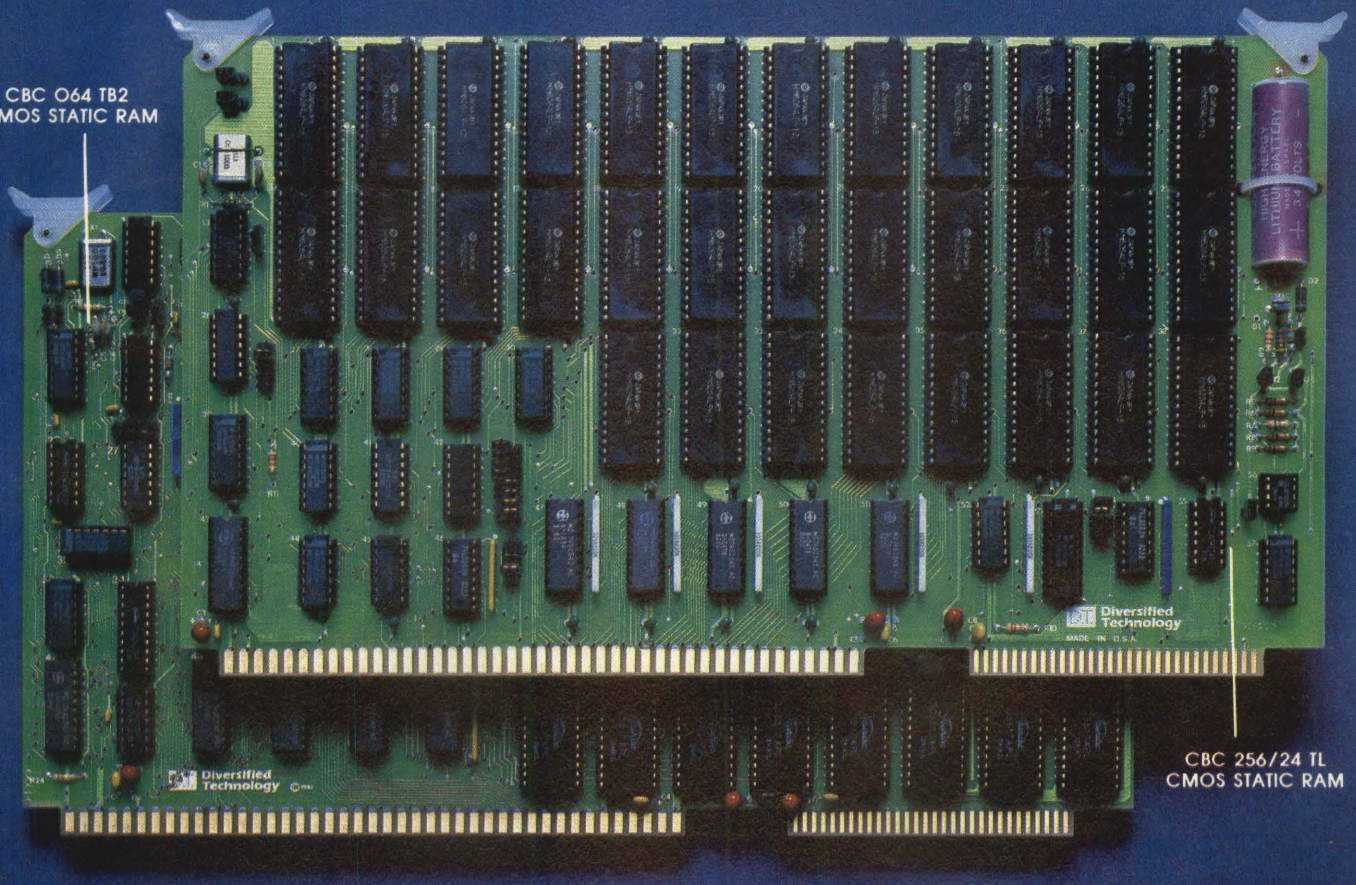
Digital system designers neglect a powerful tool for dealing with emi/rfi when they design without considering connectors. The widespread noncompliance of Class B systems with FCC docket #20780 bears witness to the consequences of such neglect. In retrofitting for compliance with FCC standards, connectors can mean the difference between successful emission control and continued violation. Connectors selected with emi in mind must cover radiated emission as efficiently as carefully developed cable and case shielding. Furthermore, the connectors must control conduction—otherwise, shielding effectiveness is diminished. The proper mix of shielded and filtered connectors can enhance a digital design by suppressing the radiation and conduction of internally generated rf interference. This kind of solution prones reliable, cost efficient and, most importantly, quiet.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 701 *Average 702* *Low 703*

NONVOLATILE CMOS RAM BOARDS

Better Performance than Bubble - at a Comparable Price



Compare these Key Features:

	INTEL iSBC* 254 - 2A BUBBLE MEMORY BOARD	DTI CBC 256/24 CMOS STATIC RAM BOARD
Bus	Multibus*	Multibus*
Memory Size	256K bytes	256K bytes
Operating Voltages	5 V, 12 V	5 V
Operating Currents	3.0A, 1.4A (max.)	100mA (max.)
Cycle Time	48 milliseconds avg.	300 nanoseconds typ.
Card Slots Required	2	1
Operating Temperature	0°-55°C	0°-70°C

ADDITIONAL FEATURES OF DTI'S CBC 256 INCLUDE:

- All - CMOS technology.
- Flexible addressing options: 16 bit with on-board bank select or 20/24 bit contiguous.
- On-board automatic memory protect.
- 8 or 16 bit data words.
- 3-year cumulative data retention time.
- 512K, 256K, 128K, 64K, 48K, 32K and 16K byte versions.

CMOS RAM NOW AVAILABLE FOR VME AND LSI-11 SYSTEMS, TOO!
CALL US FOR QUOTES ON CUSTOM RAM FOR YOUR MICROCOMPUTER BUS

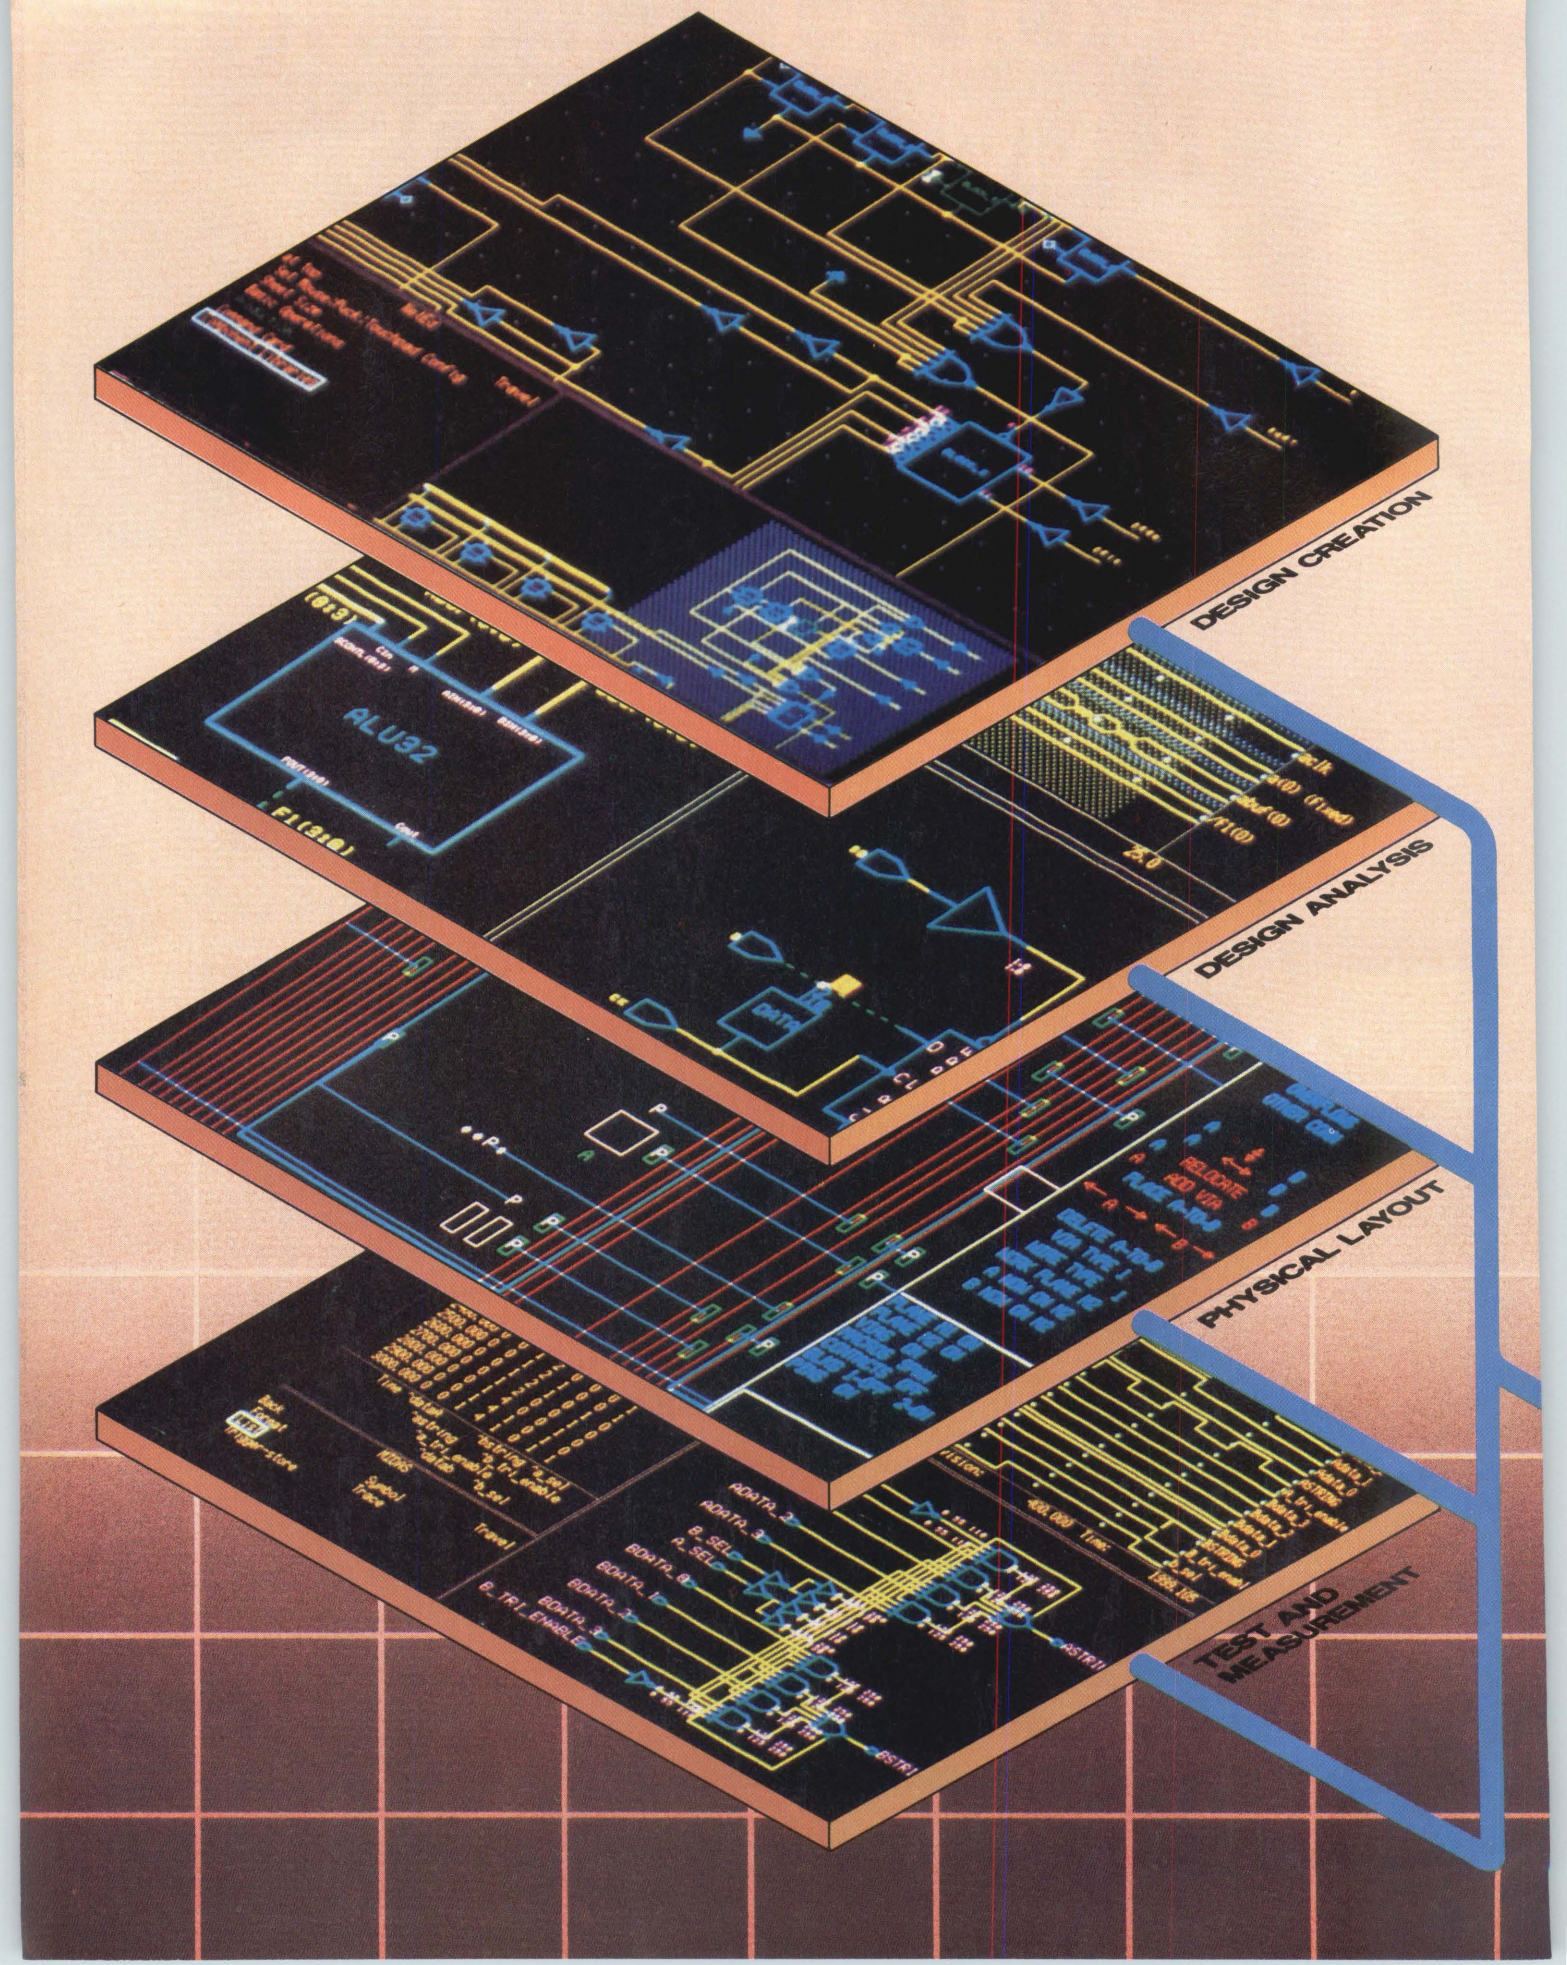
Multibus and iSBC are trademarks of Intel Corp. Above specifications taken from manufacturers current published data.

For more information regarding the CBC CMOS RAM boards, or any of our other all-CMOS MULTIBUS* boards, call or write Bill Long, CBC Product Manager at **(601) 856-4121**.

Diversified Technology
An Ergon Co.

P. O. Box 748, Ridgeland, MS 39157
Telex 585326.

CIRCLE 50


DESIGN CREATION
DESIGN ANALYSIS
PHYSICAL LAYOUT
TEST AND MEASUREMENT
ALU32

Your Mentor now makes design engineering four steps faster.

The concept of Computer-Aided Engineering has now reached its logical and highly productive conclusion:

A non-stop path from schematic entry to debugged prototype hardware. One that anticipates all your engineering needs and has the right tools waiting every step of the way.

The Mentor Graphics IDEA 1000.

Capture schematics and creativity as well.

With Mentor's powerful graphics-driven interface, the transition from concept to symbolic circuitry has never been faster.

And in addition to "flat" schematics, the IDEA 1000 lets you create an entire hierarchy of design data. From function diagrams down to transistors, you have a better conceptual grasp of your design.

Save time and money through simulation.

Mentor's digital and analog circuit simulators let you bypass much of the expense and labor associated with breadboard prototype circuitry.

These simulator tools simply access the software version of your design which resides in the IDEA 1000 system database. You head off most hardware problems before they're even physically realized.

Automate physical layout tasks.

When you're ready to take your design to physical layout, Mentor's integrated tool set tracks right along with you.

Our CADISYS gate array layout tools deliver true state-of-the-art performance. The entire gate array layout process can be completely automated from start to finish.

Use integrated logic analysis.

Mentor Graphics completes the hardware design cycle through MIDAS 7000, a fully integrated logic analysis system.

The same Mentor workstation that helped you produce your hardware will now help you verify its functionality. You can even compare real-time data acquisitions with earlier simulation runs.

Your Mentor puts it all together.

The IDEA 1000's computerized and integrated design environment is the key to faster, better electronic engineering. Contact us and we'll show you why.

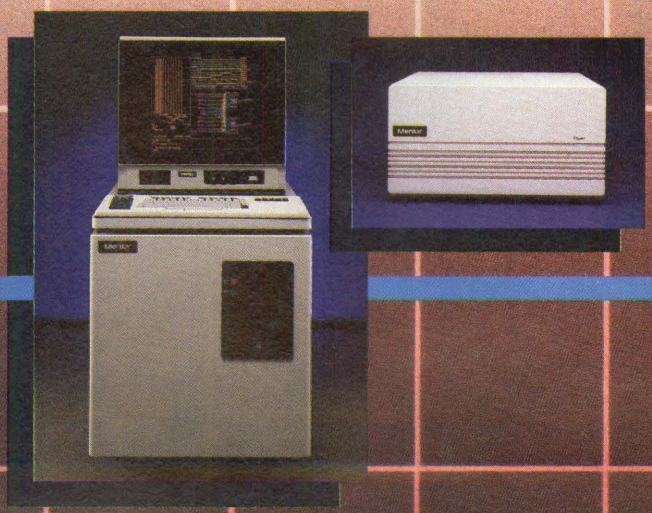
Mentor

Mentor Graphics Corporation
8500 S.W. Creekside Place
Beaverton, OR 97005
(503) 626-7000

Mentor Graphics (U.K.) Ltd.
Phone: 0734-884888

Mentor Graphics (Deutschland) GmbH
Phone: 089/319-1003

Mentor Graphics Japan Co., Ltd.
Phone: (03) 989-7950



CIRCLE 51

Gould...Innovation and Quality in UNIX-based Systems

The Firebreathers from Gould blast the competition into oblivion.



*UNIX is a trademark of AT&T Bell Laboratories
™PowerNode and PowerSeries are trademarks of Gould Inc.
™VAX is a trademark of Digital Equipment Corp.

With blazing performance.

Great leaps in raw performance are rare in the computer world. Usually, changes occur in increments of half-a-MIP or so. Now real Firebreathers roar into the arena, quadrupling the best the competition has to offer.

With scorching speed.

These creatures don't run a little faster. Running real production code supplied by VAX™ users, our PowerNode™ 9000 benchmarks at 4.5 times faster than the VAX 11/780...at a comparable price! Even our second-in-line PowerNode 6000 runs 1.5 times faster. But with all this power and speed, the PN6000 has a dainty footprint...60% less than the VAX 11/780.

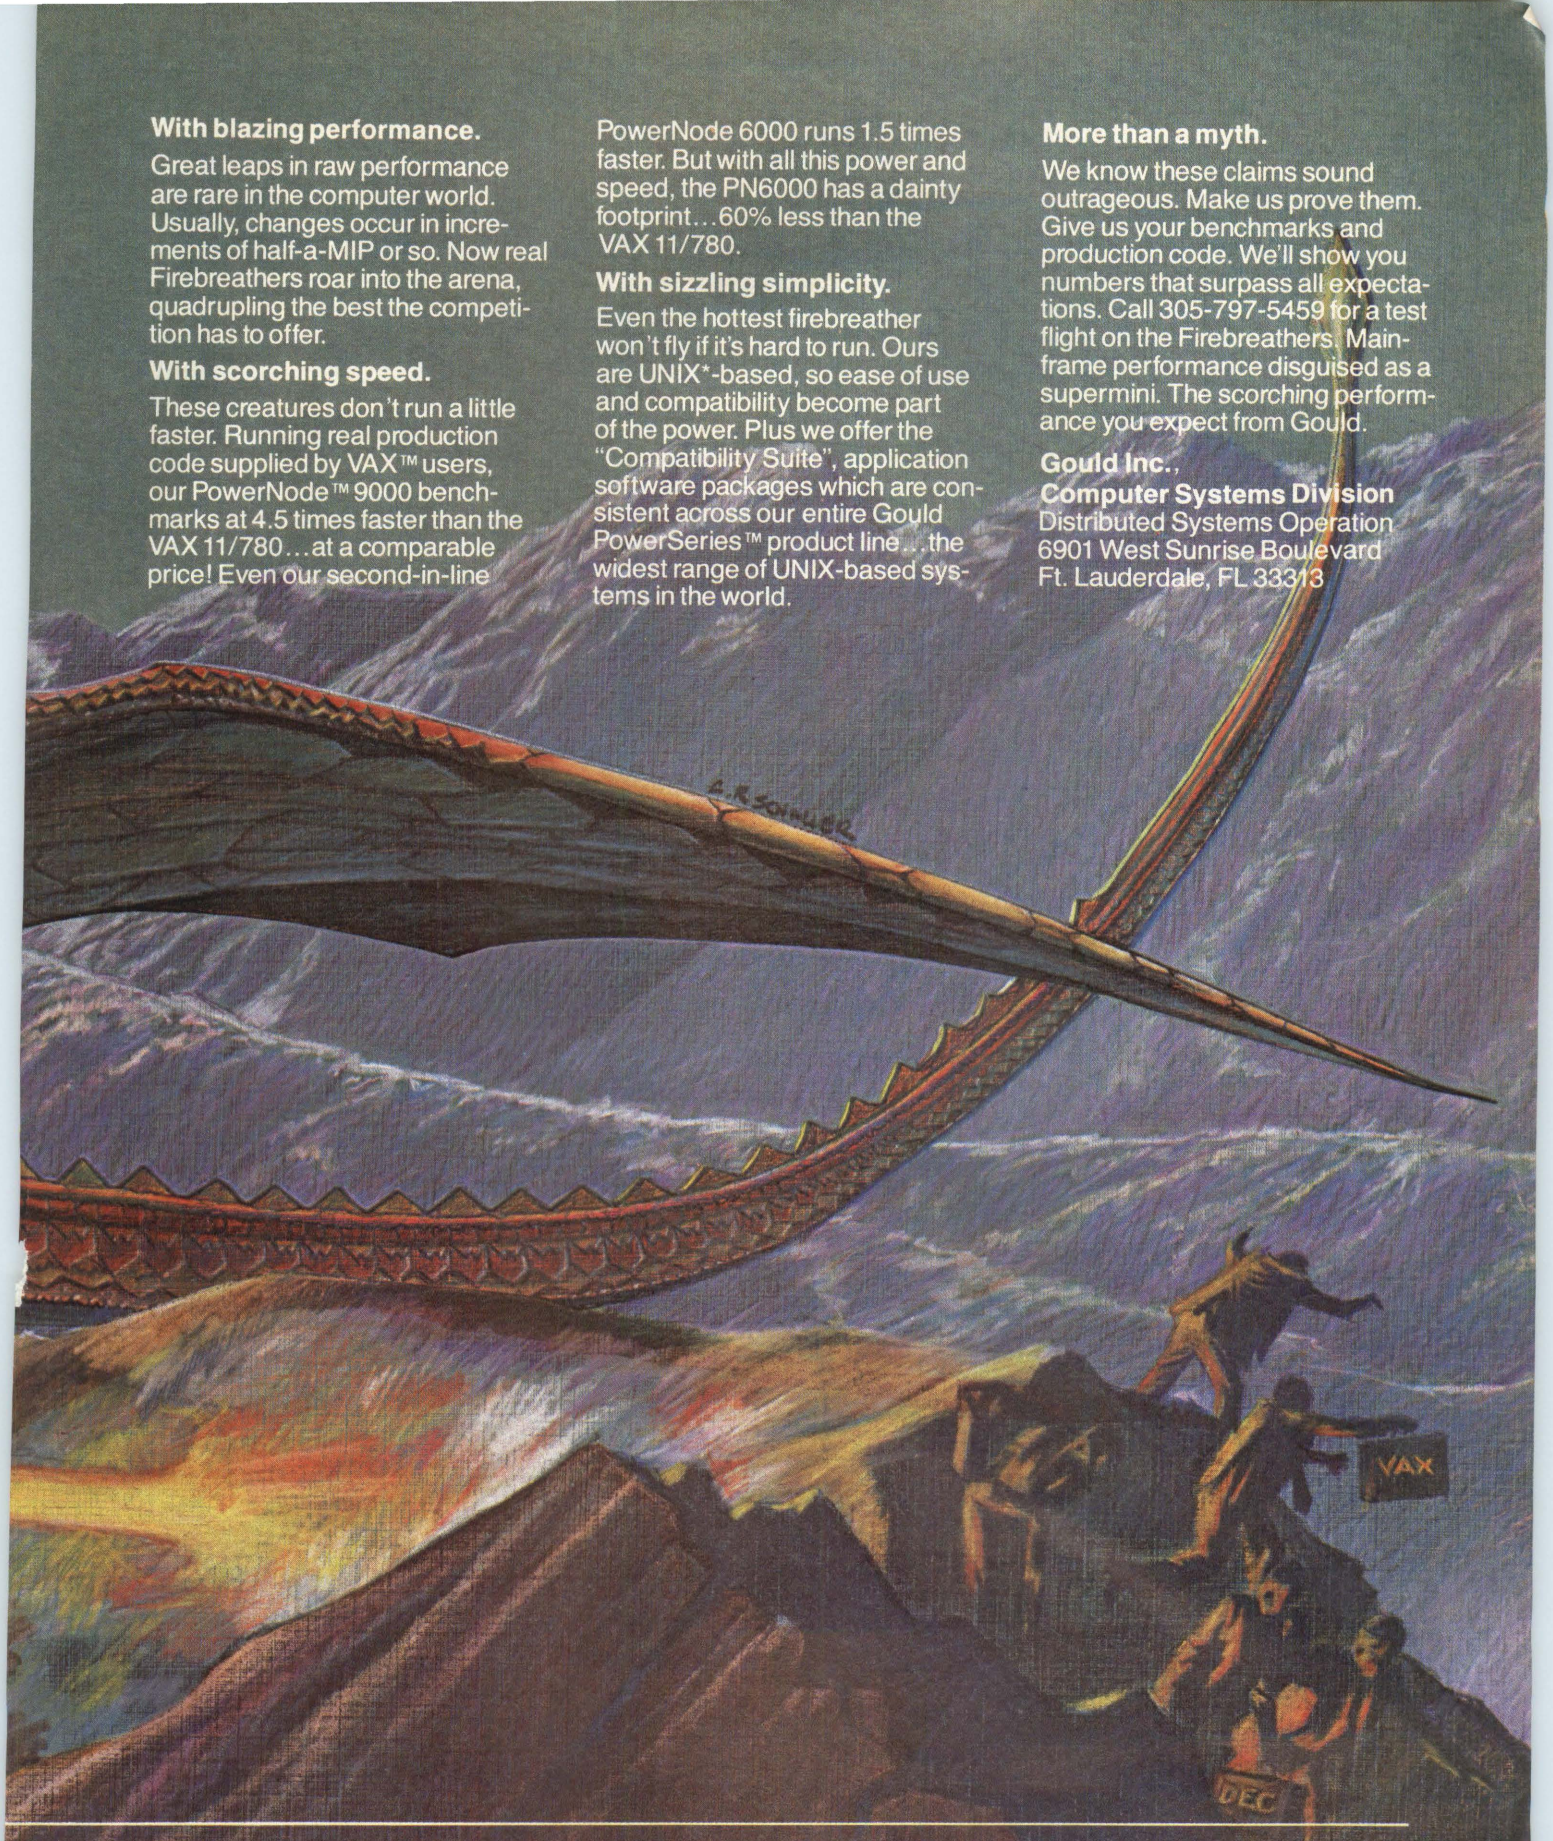
With sizzling simplicity.

Even the hottest firebreather won't fly if it's hard to run. Ours are UNIX*-based, so ease of use and compatibility become part of the power. Plus we offer the "Compatibility Suite", application software packages which are consistent across our entire Gould PowerSeries™ product line...the widest range of UNIX-based systems in the world.

More than a myth.

We know these claims sound outrageous. Make us prove them. Give us your benchmarks and production code. We'll show you numbers that surpass all expectations. Call 305-797-5459 for a test flight on the Firebreathers. Mainframe performance disguised as a supermini. The scorching performance you expect from Gould.

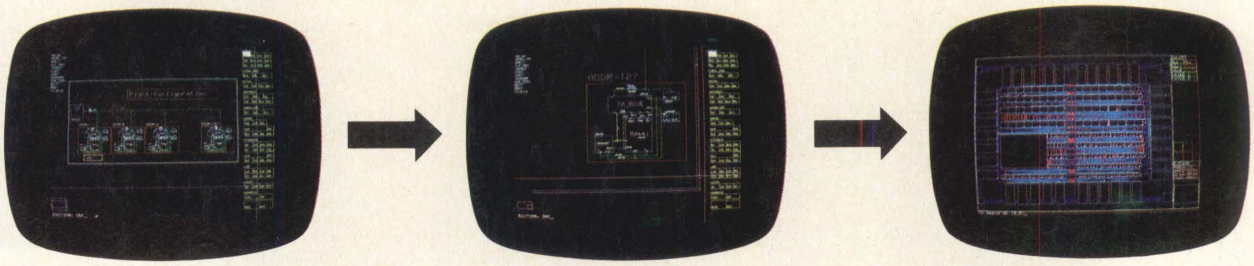
Gould Inc.,
Computer Systems Division
Distributed Systems Operation
6901 West Sunrise Boulevard
Ft. Lauderdale, FL 33313



GOULD
Electronics

CIRCLE 52

Silvar-Lisco presents hierarchical design–your way.

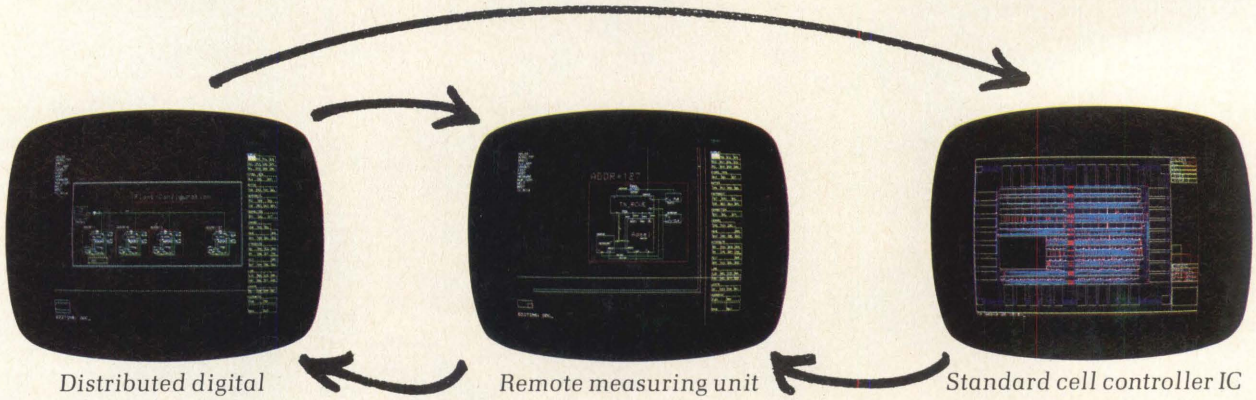


Distributed digital control system

Remote measuring unit

Standard cell controller IC

Or your way.



Distributed digital control system

Remote measuring unit

Standard cell controller IC

Hierarchical design

Our SL-2000 integrated CAE system lets you design hierarchically *and* maintain continuous design verification at every level.

Say your task is to design a control system for a chemical factory. You create a hierarchical database using our Structured Design System (SDS) for design capture. Then you use our HELIX behavioral simulator for initial system simulation. Then you partition the system and design at the technology-specific level. For example, use our STARGATE™ turnkey design system for gate arrays, or STARCELL™ for standard cell devices.

Design decisions at one level are automatically incorporated at all other levels, so you maintain integrity from concept through implementation in silicon.

Portability

Our CAE system runs on IBM, DEC/VAX, and Prime computers, *as well as* Apollo-based workstations. So you can continue using your current computers—saving the cost of hardware some other CAE systems force you to buy—or add to any existing hardware with interactive workstations supplied by us.

For more information, call or write Silvar-Lisco, 1080 Marsh Road, Menlo Park, CA 94025, (415) 324-0700. In Europe, contact Silvar-Lisco, Abdijstraat 34, B-3030 Leuven, Belgium, (016) 20 00 16, Telex 22128. In Japan, contact C. Itoh & Co. Ltd., 5-1 2-chome, Kita Aoyama Minato-ku, Tokyo 107, Japan, 03 497 3203, Telex J22295.

SILVAR LISCO™

© 1984 Silvar-Lisco

CIRCLE 53

SOFTWARE QUALITY: DESIGN IT IN FROM THE START

A comprehensive software quality assurance program should be an integral part of the development process. This program ensures not only that functions work as planned, but that planned functions are useful.

by M. Ghiassi

Developing reliable software requires a structured quality assurance program from the very beginning of the software design. Because the programmer/developer tends to be biased in favor of the product, and is under pressure to meet product development deadlines, the customer is often the first real tester of new microprocessor software. If software quality assurance is not a direct part of the program development, a healthy exchange of ideas during testing is eliminated from the development stage. As microprocessor systems approach the mainframe performance level, and system software as well as application software becomes much more sophisticated, better software quality assurance methods become indispensable.

Design methodologies, test criteria, and specialized tools contribute to a successful software quality assurance (SQA) program, as well as the experience level of the quality assurance (QA) personnel. Independence from the design group is a must, as the QA group must also act as "first customer" in evaluating a new software product.

According to Barry Boehm of TRW (El Segundo, Calif) and Thomas Standish of the University of California at Irvine, the national demand for software engineers is increasing by 12 percent annually, while the available supply is increasing by only four percent a year.[1] Even with this four-percent increase in individual programmer productivity, there is still a four-percent increase in the gap between supply and demand. By the 1990s, a shortage of over a million programmers is expected.

There are two ways to close this gap: by increasing productivity even further, and/or by cutting down the time spent on maintenance, which is the number one software activity. Approximately 60 to 80 percent of software cost and 70 percent of programmer time are spent in maintaining existing programs. Obviously, time spent maintaining old programs detracts from time available for the creation of new programs demanded by the industry. Many industry insiders estimate that, while it takes $40 to $60 to write a line of code, to maintain it over a period of time can cost as much as $4000.

Meeting the need for SQA programs

Therefore, the time and cost required for program maintenance, with an expected cost increase in the future, give rise to a need for effective SQA programs. The use of SQA can cut down the time spent on program maintenance by uncovering and eliminating potential problems before the product leaves the development lab. In addition, putting an SQA program in place throughout the software product development cycle cuts down on the time spent later in the costly and time-consuming process of program maintenance.

Typically, there has been no separate SQA function in software development departments; it has been common for programmers to perform their own QA evaluation, and to develop their own test tools. Invariably, a busy development department

M. Ghiassi is a member of the technical staff at National Semiconductor, 2900 Semiconductor Dr, Santa Clara, CA 95051. Dr Ghiassi is also assistant professor of information systems at the University of Santa Clara. He holds a BS in economics from the University of Tehran, Tehran, Iran, an MS in economics from Southern Illinois University, and an MS in computer science and a PhD in engineering, both from the University of Illinois at Urbana.

will underestimate and underbudget the resources required for testing and QA.

In addition, programmers/developers suffer from the bias of the work environment. Often they are so familiar with the program and how it "should" operate that it is difficult to devise valid tests for unexpected glitches. Even very good programmers do not necessarily know how to test due to lack of experience. As a result, the customer often becomes the first "critical user," and by default, the QA backstop. Once a program is released to the public, it becomes difficult, if not impossible, to halt its momentum long enough to adequately resolve all the bugs uncovered in the field.

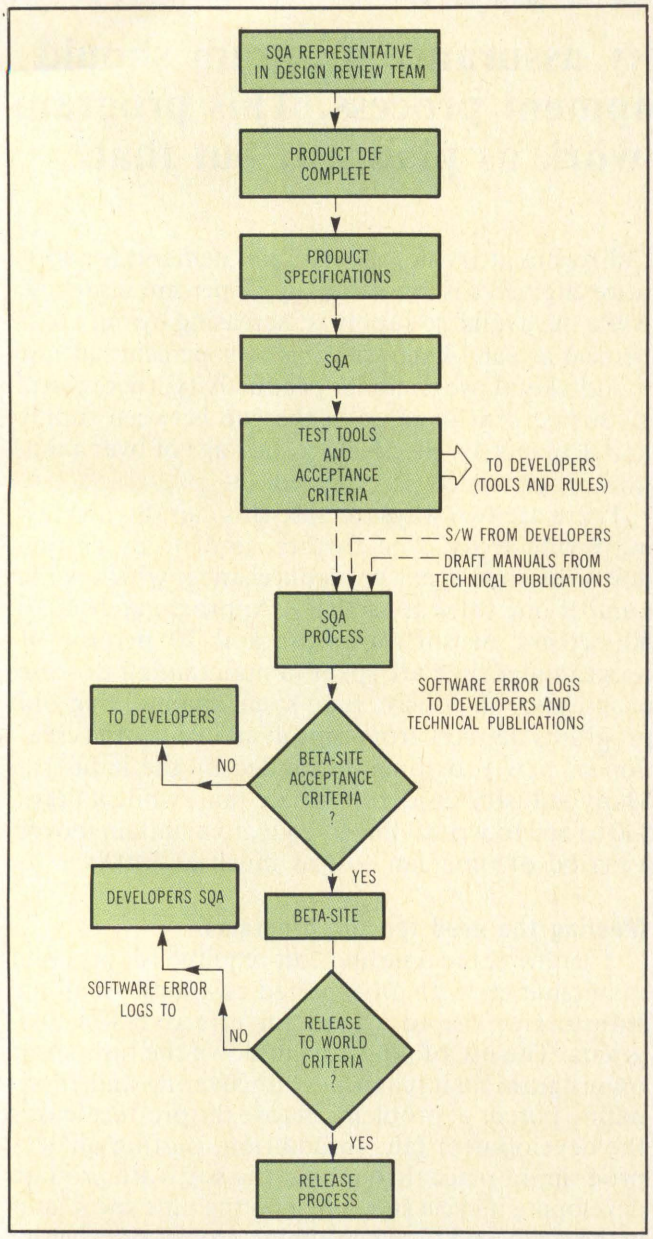


Fig 1 In product definition, the software quality assurance (SQA) group advises the development team about potential problems that can arise during testing, defines the acceptance criteria, and learns about the product in advance of final testing during the release stage. Once the software has passed SQA criteria, it is released to beta-sites, and if all tests are passed, it is released to manufacturing.

To avoid this situation, the SQA group should be the first "critical user," apart from the software development team. This eliminates the bias of the programmer/tester and prevents critical omissions due to overfamiliarity with the product. Furthermore, SQA sets standards for software operation. A set of comprehensive standards ensures that the software product will be able to stand the test of time and field trials with actual customers. Therefore, tests devised by the SQA group should set up much more stringent conditions than those which will be encountered by users.

Testing and quality assurance

Although the goals of testing and QA appear to be the same, there are important differences. Testing is done during the development phase and determines whether or not the product meets the defined standards; QA is oriented toward defining what those product standards should be. Unlike testing, which asks whether or not the product performs according to what the designer has in mind, QA helps decide if the product does what the customer/user needs, or expects it to do. Pfan defines QA as "the name given to activities performed in conjunction with the development of a software product to guarantee the product meets the specific standards. These activities reduce doubts and risks about the performance of the product in the target environment."[2]

Thus, QA activities should begin at the onset of the design cycle, and continue throughout the product development, testing, and manufacturing processes. In the past, software product development without a reliable QA program has resulted in functionally correct, but inadequate software that, although meeting the designer's or design group's specifications, fell short of meeting the needs of the intended user group. Unfortunately, much microcomputer software falls into this category.

The establishment of an independent SQA team within the framework of a software development group provides an independent, unbiased viewpoint from which to judge the functionality of a piece of software. In other words, it furnishes a specialized resource—ie, experienced personnel who can deal with the problems that frequently occur in software development. In the SQA group at National Semiconductor, for example, there are operating system experts, compiler testers, and utility specialists who have accumulated years of experience in testing these different software categories.

As shown in Fig 1, QA should be part of the early stages of design review. Proper software development requires the cooperation of many groups, including software engineering, marketing, technical



Now you can watch the entire Whizzard® family in 3D.

Because the industry's widest range of high performance graphics systems now has a whole new dimension.

A third dimension.

Which means now you can buy 3D Megateknology starting for *under* $12,000. For instance, the Whizzard 1600 desktop design terminals with color or monochrome capabilities, very high resolution and VT-100™ compatibility.

Or our Whizzard 3355, which holds the speed record for its price range—400,000 vectors per second—thanks to our one-of-a-kind Graphics Engine.™

Then there's the ultimate computer graphics system: the Whizzard 7200. It guarantees remarkable speed and flexibility, modular architecture that supports high resolution and real-time dynamic color raster displays, and unmatched interactive graphics capabilities.

Plus, with every one of our products you get Megateknology—our longstanding trademark of innovative design, quality and high reliability.

There's a whole new dimension to our Whizzard family, thanks to the latest in 3D Megateknology.

Watch it add more dimension to your creativity.

MEGATEK CORPORATION®
A UNITED TELECOM COMPANY
Making History out of State-of-the-Art

That's Megateknology.™

World Headquarters • 9645 Scranton Road • San Diego, California 92121 • 619/455-5590 • TWX: 910-337-1270
International Headquarters • 50, Avenue Des Boveresses • CH-1010 Lausanne, Switzerland • Telephone: 41/21/335016 • Telex: 25037 mega ch

VT-100 is a trademark of Digital Equipment Corporation.

CIRCLE 54

```
main( )

  int def;
  double wrongtype;

  def = 0;
  wrongtype = 3.0;
  switch(wrongtype)

    case 3:
      def++;
      printf("FAIL: An expression of type 'double' allowed n");
      printf("       in a SWITCH statement. Note that the n");
      printf("       statements for 'int' 3 were executed. n");
      break;

    default:
      def++;
      printf("FAIL: An expression of type 'double' allowed n");
      printf("       in a SWITCH statement. Note that the n");
      printf("       'default' statements were executed. n");

  if(def==0)

      printf("FAIL: An expression of type 'double' allowed n");
      printf("       in a SWITCH statement. n");
```

Fig 2 The test program determines whether or not the C compiler issues an error message for a switch statement whose expression has a result type other than 'int' (integer), after the usual arithmetic conversions have been performed.

publications, hardware engineering, and SQA. These groups comprise the software design review team. The responsibilities of the SQA group include being aware of what the product is about and raising developer awareness to the need for SQA.

If the product is new, the process must allow time for SQA to learn and grow with the product, from the specification stage through the many phases of product development. With such early involvement, schedules can be planned, and procedures for testing the product can be put in place before the crunch of final testing occurs. Development of new tools, if necessary, testing of the product against established benchmarks, and setting up and testing recommended hardware configurations all fall within the jurisdiction of the SQA group.

The literature on software maintenance points out that the later in the development process that a bug is found, the more difficult and expensive it is to correct. Thus, SQA should alert the development team to the importance of early testing. SQA provides an awareness of test cases and the availability of test suites and tools for developer use. (Individual test programs are used to check particular functions or parameters of a product, while a collection of test programs that exhaustively tests an entire software product is called a test suite.) The SQA group must create and maintain a database history of test cases against which a product can be measured. They should also provide a set of test suites and tools that can be used for systematic analysis of a software product throughout its life cycle.

Examples of test suites are those used by the U.S. Department of Defense (DoD), National Bureau of Standards, and Western Electric's Bell Labs to test conformance of compilers (Ada, Cobol, and Fortran), and operating systems (Unix). In the case of Ada, for example, the DoD has developed a rigorous testing procedure to certify the proper operation of Ada compilers. It has also certified some third-party software houses to test commercial compilers. Similarly, AT&T has set up ground rules for an acceptable System V for its 16/32-bit microprocessors, to pass these tests in accordance with its predefined rules and procedures. Typically, the software development team rarely takes the time to do these types of tests, and this testing becomes part of the independent SQA group charter.

Looking at the SQA procedure

Once a piece of code has been developed by the design team, it is not automatically ready to be transferred to the SQA group. Release of software to SQA should be accompanied by the following: source code files (object files as well, when appropriate); a complete list of files for each component; a complete set of manuals or, at least, specifications; installation procedures; documents describing test procedures and the test suites used by developers with a list of known bugs; all supporting equipment, when necessary; and specific information on expected customer configurations.

After participating in the design review sessions and approving the design team testing, SQA is ready to test the product fully. This task requires testing the following categories: conformity to specifications; error handling; boundary conditions; deviations from standards [eg, International Standards Organization (ISO) Pascal specs]; and extensions (definitions and exceptions).

In the first case, there are various checks that must be made to ensure that the product does indeed meet the stated specifications as defined by the initial product design team. For example, if the product is a language, its grammar and syntax are defined in the product specification. SQA must establish the correctness of the spec by creating test cases for checking each type of statement and combination of language elements.

Fig 2 gives an example of a test program designed to check a C compiler feature. In this case, the test program deliberately incorporates an error, followed by a print statement that denotes the error. If the compiler correctly flags the error, the statement will not be printed. If it does not detect the error, the program will be compiled, executed, and the fail message printed. Test programs such as this one are developed to test programming language features—eg, to test grammar, parameter, and argument lists, proper variable names, etc.

Proper error handling is tested by feeding the compiler incorrect constructs and seeing if the compiler handles them correctly. Whereas the conformity-to-spec testing checks to see if the program handles correct data properly, error handling tests confirm that errors are correctly detected, reported and, when possible, corrected.

Importance of proper error handling

An example of the problems that can crop up when proper error handling tests are not available is apparent in a popular database management program for personal computers. In this program, users can create files with file names that are considered illegal by the operating system. This can lead to the creation of files filled with important data which cannot be accessed. Since the operating system does not recognize the illegal file names, it will not permit the opening, or even renaming of the files. For this reason, ensuring proper error handling as opposed to simply confirming correct function implementation is vital to database management.

Often a specific language implementation, such as ISO Pascal, must take into account the details of the underlying hardware system. Thus, some features may perform differently on different hardware. In such cases, documentation for every deviation from the expected, as well as any alternative method of achieving the desired result, should be written.

Some examples of deviations from Pascal standards are the specifics of I/O handling and pointer size. If the product does not conform to standard usage, the software development team should document each deviation and include methods of working around the deviation. SQA's job is to see that this documentation is actually there, and that the alternative methods outlined in the documentation do indeed work. The same is true of extensions to a language. All manufacturer-specific enhancements must be documented and integrated so that they do not interfere with the standard language subset.

Boundary conditions are checked by determining the legal values of expressions or parameters, and running programs that exceed those values. An example of such a test program is shown in Fig 3, a Pascal program to check for an illegal value in an expression involving a MOD operator. This program uses the Pascal MOD operator with an illegal value. The left-hand term of an expression using the MOD operator cannot have a value of 0 or less. In this case, a negative value is used and a fail message is printed if the program is compiled. If the error is detected, it will be flagged as a runtime error. If not, the program will run, and print the error message.

It is essential that boundary conditions be checked for all parameters and conditions. In the QA process for an operating system, some examples of boundaries are the maximum limit of active users, and the maximum number of active processes per user. The SQA's function is to try to break all the rules, and to see if the software product can detect the error, report the error, and operate around the error in the intended manner.

```
Program NegMOD (OUTPUT);

Var
  I, J, K: Integer;

Begin
  I:= 3;
  J:= -2;
  K:= I Mod J; (*This is a run-time error *)
  Writeln('FAIL: I MOD J not flagged when J    0');
End.
```

Fig 3 The Pascal test program determines whether or not a runtime error message is issued for an illegal left-hand term of an expression involving the MOD operator. The left-hand term of such an expression is not allowed to have a value of 0 or less. The test program is concerned with what happens when the left-hand term has a value of less than 0.

An example of a standard test suite for the test of a language such as ISO Pascal is the Tasmanian test suite developed at the University of Tasmania, Australia. The Tasmanian test suite incorporates sections that test for spec conformance, language deviation, error handling, implementation-dependent characteristics, measurements of code densities, and operational speeds.

During the last several years, the use of such standardized test suites has become the rule for many SQA groups concerned with system software. As previously mentioned, there are standard test suites for Ada, Fortran, Cobol, ISO Pascal and the Unix operating system. Unfortunately, there are no such standards available for application programs. Thus, the presence of a strong SQA group is even more important for software developers aiming at the more diverse application market.

Regression tests and software error logs

The SQA group's final responsibility is that of providing a data history repository for the software product, from inception to final release. Since revisions are inevitable, SQA must fully document all changes to the product over time. It should also confirm the proper operation of the software as these changes are made. One useful method of ensuring proper operation during such changes is regression testing.

Software products usually go through several modifications. In addition, construction of test suites allows for automation of the testing process. After each modification, the system must be retested through regression testing. One major concern in regression testing is the level of testing that is required when only a portion of a system is changed. If the modification of the system can be isolated

safely from other parts of the software, then local testing of that software portion may be sufficient.

An important point to remember, however, is that a correction or modification in one section of the system can impact other portions that have not been modified. This danger is always present and local testing introduces a certain degree of risk to the regression testing process. If the results of the regression testing are the same, the software can be said to be functionally unchanged. These types of tests are particularly useful in confirming proper operation of systems involving modifications of both hardware and software.

Making changes

Another useful tool in documenting software changes is the software error log (SEL). The SEL should report the history of software bugs from their first occurrence, suggest workaround solutions, and document the corrected resolution of the problem. There are three categories of information in the SELs used at National Semiconductor: environmental information, error definition, and error tracking.

SEL environmental information points out the system where the error occurred, and describes the SEL ID, the software component name, the host operating system, the host hardware configuration, and the severity of the error. Error definition, the second category, is a description of the symptoms associated with error, documentation references, test case references, expected results of error correction, and workaround solutions to the error for use with the existing revision. Finally, error tracking, or history of the bug, logs the date first encountered, other versions of the software that also have the error, bug verification by other users, and any other fixes or experiences by other team members who have found the error.

In addition to providing a continuous history of the software product, the SEL report provides a structured methodology for determining where valuable resources should be spent in improving the program. The severity rating of the SELs along with with number of errors can be a determining factor in ordering priorities that will define which errors should receive the most attention during the revision stages. A total count of SELs for each software component is also used as a guideline for whether or not the product is releasable. At National, there are three release levels with different sets of acceptance criteria: release from the development group to SQA; release to beta-site installations; and release to manufacturing.

The development group provides a complete set of source files, documentation for constructing customer configurations of the product, a description of the testing procedures that were used in engineering with a summary of test results, a description of known errors or deviations from original specifications, and a formal release document. The latter includes the installation procedures with verifying routines, and a description of the development environment. At the same time, the SQA group should receive the latest copy of the manuals from technical publications, usually in draft form.

At this point, SQA starts comprehensive testing of the product. Test cases are employed, test suites run, and all rules are broken in search of undocumented phenomena and incorrect error handling. The SQA group uses the SELs to find and classify all errors into high, medium, and low degrees of severity. Typically, high severity errors affect product usefulness, medium ones cause unexpected results of undocumented features in the product operation, and low severity errors cause minor inconvenience.

Before the product can be released to a beta-site, all high severity errors must be resolved and a report summarizing software behavior when medium and low severity errors are encountered should be written. Accompanying this report should be a suggested workaround solution to each problem. In addition, when released to a beta-site, the software should be accompanied by a complete set of manuals necessary for installation and operation. And finally, a designated customer support contact should be available. For release to manufacturing, all SELs should be resolved, a positive report from the beta-site should be confirmed by SQA personnel, and final customer support facilities should be in place.

When an SQA group is an integral part of the software development cycle, the ability to produce consistently reliable software products is improved. Using SQA techniques throughout the development cycle cuts down the time required for later program maintenance. This also ensures that the product is released according to the agreed-upon specifications with all deviations and revisions documented and classified.

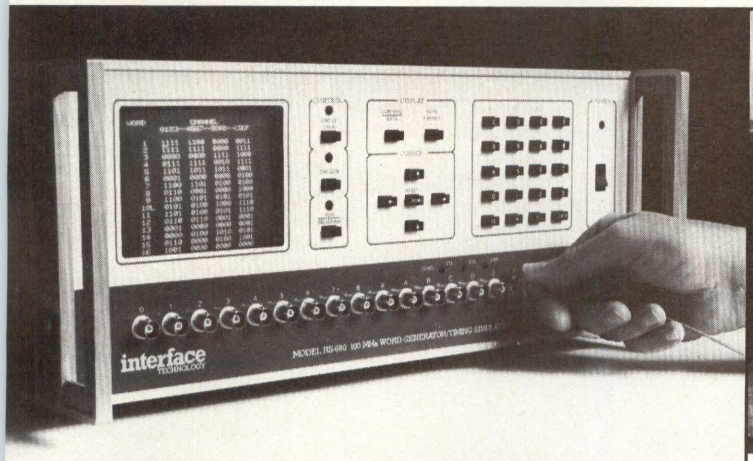
References

1. B. W. Boehm, and T. A. Standish, "Software Technology in the 1990s: Using an Evolutionary Paradigm," *Computer* (IEEE), Nov 1983, pp 30-37.
2. P. Pfan, "Applied Quality Assurance Methodology," *Proceedings of the Software Quality Assurance Workshop,* Association for Computing Machinery, Nov 1978, pp 1-8.

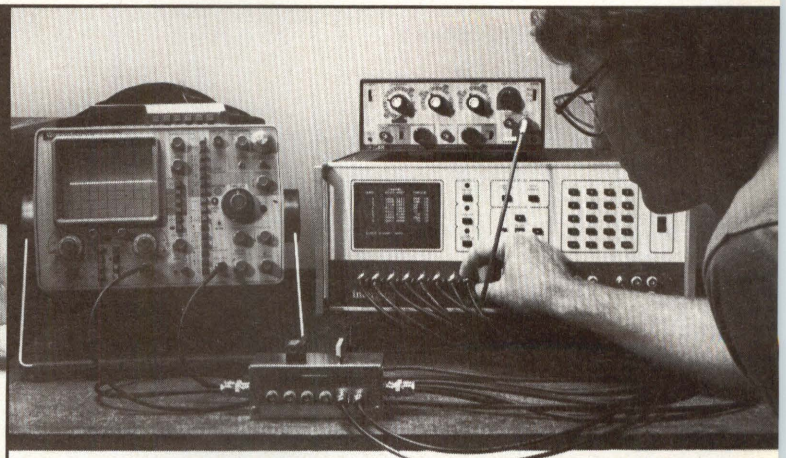
Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 704 *Average 705* *Low 706*

WHO'S GOT DIBS ON THE DWG?

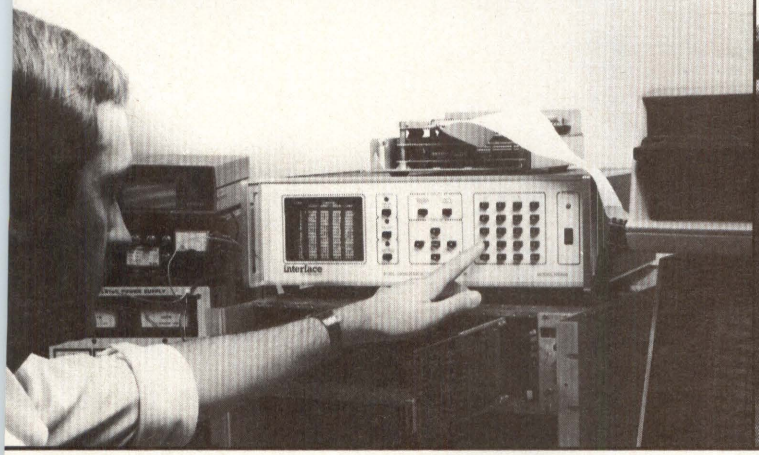


Digital Word Generators are always in demand.

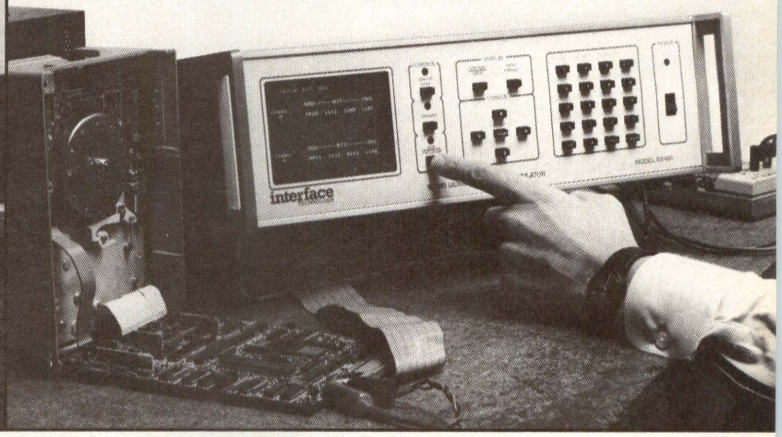


DWGs speed up IC characterization.

DWGs simulate prototype environments.



DWGs eliminate building dedicated circuitry.



You never know where a Digital Word Generator from Interface Technology will show up next.

From testing discrete designs before system integration to simulating prototype environments, these DWGs can handle it all. Speeds to 100 MHz and up to 4096 bits of memory on each of 16 channels make the RS-660 and RS-680 the most powerful and versatile instruments of their type.

With Interface Technology's unique timing simulation capability, you can generate variable width pulses or timing patterns, thus conserving time and memory space. Data may then be output in a single step, burst or continuous mode.

Both the RS-660 and RS-680 have integral CRTs. The ability to display a large portion of your serial data stream or parallel pattern makes either instrument easy to use. Use is further simplified by the choice of either hex or binary display; combined with a block memory fill for repetitive pattern entry.

Other features such as a programmable sync output for triggering, remote interface, and up to a 64K deep serial memory just add to an already powerful instrument.

So put an Interface Technology RS-660 or RS-680 DWG to work at your place—all over your place. Call today to put "dibs" on your DWG. (714) 592-2971.

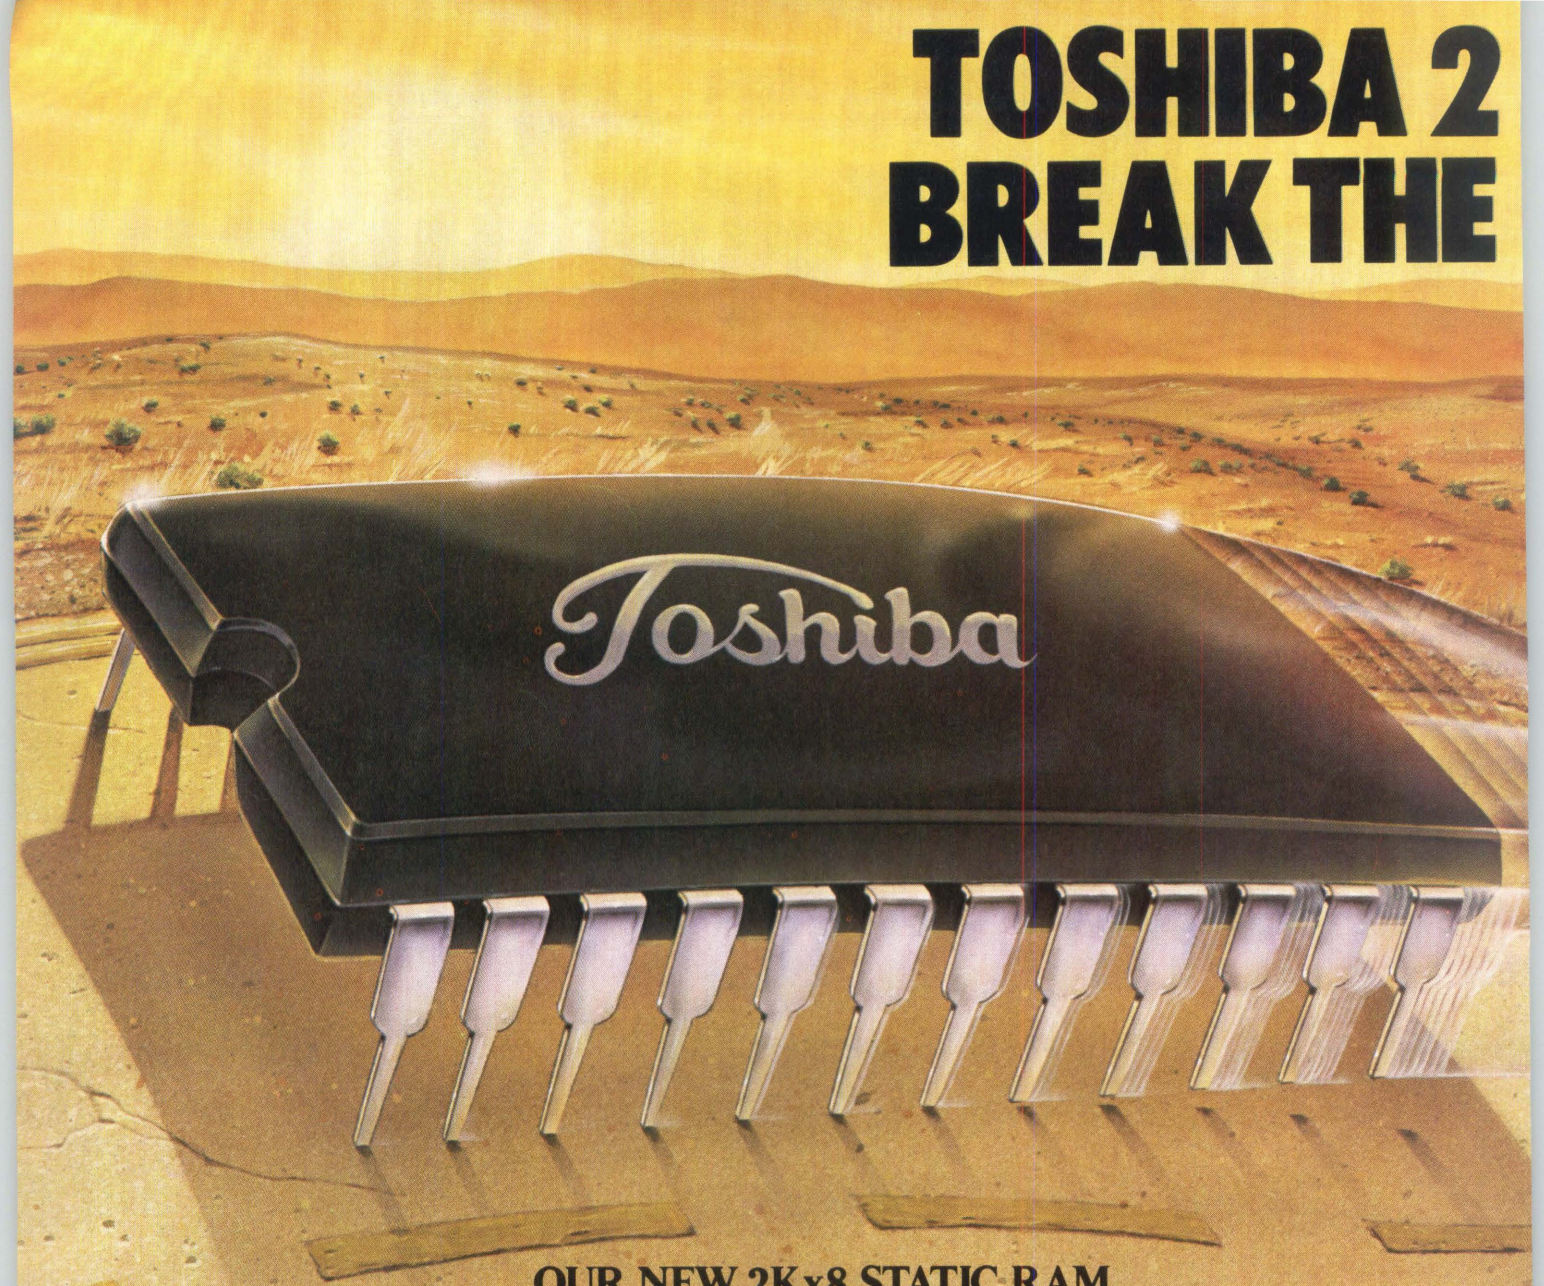
interface TECHNOLOGY

If it's digital, we can test it.

A Dynatech Company, 150 East Arrow Highway, San Dimas, CA 91773

CIRCLE 56

TOSHIBA 2 BREAK THE



OUR NEW 2Kx8 STATIC RAM IS TWICE AS FAST AS ANY OTHER BYTE-WIDE. HITTING SPEEDS TO 45ns.

Toshiba has the world's fastest 2K x 8 Static RAM. With speeds as fast as 45ns and other byte-wide units with power consumption as low as 1μA, your range of design options just got twice as wide as before.

Our new TMM2018D provides both high-speed and low-power features with an access time of 45ns. This, along with high density, explains why they're rapidly displacing bipolar devices.

All our high-speed *NMOS* and *CMOS* 2K x 8's are designed for maximum compatibility with microprocessor bus structures.

In fact, ours were the first 16K CMOS RAMs on the market. We designed them for a maximum 1μA standby current.

Operating from a single 5V power supply, our byte-wide RAMs are available in a 24-pin package, DIP (.300″ or .600″), flat pack and a variety of other configurations.

AREA SALES OFFICES: WESTERN AREA, Toshiba America, Inc., (714) 752-0373; **CENTRAL AREA,** Toshiba America, Inc., (612) 831-2566; **EASTERN AREA,** Toshiba America, Inc., (617) 272-4352; **NORTHWESTERN AREA,** Toshiba America, Inc., (408) 244-4070. **REPRESENTATIVE OFFICES: ALABAMA,** Glen White Associates, (205) 883-7938; **ARIZONA,** Semper Fi Sales Company, (602) 991-4601; **ARKANSAS,** Technology Sales Company, (214) 380-0200; **CALIFORNIA** (Northern), Elrepco, Inc., (415) 962-0660; **CALIFORNIA** (Southern), Bager Electronics, Inc., (714) 957-3367; **COLORADO,** Duffy Associates, (303) 595-4244; **CONNECTICUT,** Datcom, Inc., (203) 288-7005; **DELAWARE,** Vantage Sales, (609) 663-6660; **FLORIDA,** Donato & Associates, (305) 522-2200, (305) 352-0727, (813) 785-3327; **GEORGIA,** Glen White Associates, (404) 441-1447; **IDAHO,** Components West, (206) 271-5252; **ILLINOIS,** L-TEC, (312) 773-2900; **INDIANA,** Leslie M. DeVoe Company, (317) 842-3245; **IOWA,** J.R. Sales Engineering, (319) 393-2232; **KANSAS,** R.R. Burton & Associates, (816) 763-5385; **KENTUCKY,** Leslie M. DeVoe Company, (317) 842-3245; **LOUISIANA,** Technology Sales Company, (214) 380-0200; **MAINE,** Datcom, Inc., (617) 891-4600; **MARYLAND,** Glen White Associates, (301) 252-6360; **MASSACHUSETTS,** Datcom, Inc., (617) 891-4600; **MICHIGAN,** R.C. Nordstrom & Company, (313) 559-7373; **MINNESOTA,** Quantum Sales, Inc., (612) 884-4700; **MISSISSIPPI,** Glen White Associates, (205) 883-7938; **MISSOURI,** R.R. Burton & Associates, (816) 763-5385; **MONTANA,** Components West, (206) 271-5252; **NEBRASKA,**

K x 8 RAMs SPEED LIMIT.



Part No.	Type	Access Time (Max)	Operating Current (Max)	Standby Current (Max)
2016AP-90	NMOS	90ns	80mA	7mA
2016AP-10	NMOS	100ns	65mA	7mA
2016AP-12	NMOS	120ns	65mA	7mA
2016AP-15	NMOS	150ns	65mA	7mA
*2015AP-90	NMOS	90ns	80mA	7mA
*2015AP-10	NMOS	100ns	65mA	7mA
*2015AP-12	NMOS	120ns	65mA	7mA
*2015AP-15	NMOS	150ns	65mA	7mA
*2018D-45	NMOS	45ns	120mA	20mA
*2018D-55	NMOS	55ns	120mA	20mA
**5516AP	CMOS	250ns	55mA	30 μA
**5516AP-2	CMOS	200ns	55mA	30 μA
**5516APL	CMOS	250ns	55mA	1 μA
**5516APL-2	CMOS	200ns	55mA	1 μA
**5517AP	CMOS	250ns	55mA	30 μA
**5517AP-2	CMOS	200ns	55mA	30 μA
**5517APL	CMOS	250ns	55mA	1 μA
**5517APL-2	CMOS	200ns	55mA	1 μA
5517BP-20	CMOS	200ns	25mA	30 μA
5517BPL-20	CMOS	200ns	25mA	1 μA
5518BP-20	CMOS	200ns	25mA	30 μA
5518BPL-20	CMOS	200ns	25mA	1 μA

*.300" wide dual in line package.
**Also available in a plastic flat pack (small outline package). Under development.

If you're designing cache memory, high-speed storage, hand-helds and other high-density memory applications, write for more information to Toshiba America, Inc., 2441 Michelle Drive, Tustin, CA 92680, (714) 730-5000. Or call your local distributor or sales representative.

Toshiba America broke the speed limit so there'll be fewer design limitations for you.

TOSHIBA AMERICA, INC.

A WORLD STANDARD IN MOS.

CIRCLE 57

ANDREWS

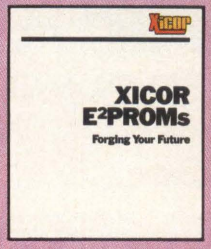
R.R. Burton & Associates, (816) 763-5385; **NEVADA,** Elrepco, Inc., (415) 962-0660; **NEW HAMPSHIRE,** Datcom, Inc., (617) 891-4600; **NEW JERSEY,** Necco 1, (201) 461-2789, Vantage Sales, (609) 663-6660; **NEW MEXICO,** Semper Fi Sales Company, (602) 991-4601; **NEW YORK,** Necco 1, (201) 461-2789, PI-tronics, (315) 455-7346; **NORTH CAROLINA/SOUTH CAROLINA,** Glen White Associates, (919) 848-1931, (615) 477-8850; **NORTH/SOUTH DAKOTA,** Quantum Sales, Inc., (612) 884-4700; **OHIO,** Del Steffen & Associates, (216) 461-8333, (419) 884-2313, (513) 293-3145; **OKLAHOMA,** Engineering Sales Company, (918) 493-1927, Technology Sales Company, (214) 380-0200; **OREGON,** Components West, (503) 684-1671; **PENNSYLVANIA,** Del Steffen & Associates, (412) 276-7366, Vantage Sales, (609) 663-6660; **RHODE ISLAND,** Datcom, Inc., (617) 891-4600; **TENNESSEE,** Glen White Associates, (615) 477-8850; **TEXAS,** Technology Sales Company, (512) 346-9940, (713) 266-2473, (214) 380-0200; **UTAH,** Duffy Associates, (303) 595-4244; **VERMONT,** Datcom, Inc., (617) 891-4600; **VIRGINIA/WEST VIRGINIA,** Glen White Associates, (804) 237-6291, (804) 295-0435, (804) 224-7764, (804) 224-0404; **WASHINGTON,** Components West, (206) 885-5880; **WISCONSIN,** L-TEC, (414) 774-1000; **WYOMING,** Duffy Associates, (303) 595-4244; **CANADA,** Electro Source Inc., (416) 675-4490.



Xicor E²PROMs. Forging Your Future.

If you're designing high-performance/high-capacity E²PROM systems, Xicor has guaranteed your future. With the right E²PROM—created for foresighted system designers.

Xicor's 64K E²PROM, the X2864A, offers the design innovations that make it the right choice. For today and tomorrow. Innovations like automatic page write for super-fast, write-time performance. And $\overline{\text{DATA}}$ polling for write verification.



All in the JEDEC-standard footprint for 28-pin memories.

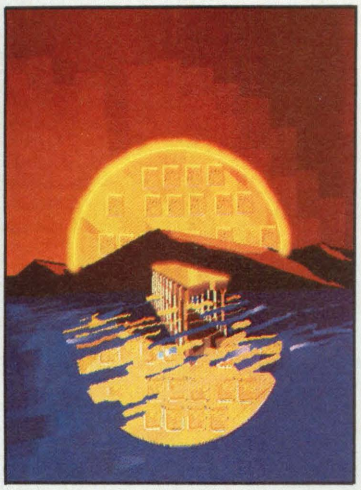
The X2864A also features an expandable architecture. Leaving pin one free to ensure your migration to 256K—without redesigning your system. Fact is, we'll be sampling our 256K E²PROMs in the first half of 1985.

That's sooner than anyone expected.

For more information, write us at 851 Buckeye Court, Milpitas, CA 95035. Or call (408) 946-6920 and ask for a copy of our free brochure.

Xicor

CIRCLE 58



Special report on semiconductor memories

103 Introduction

105 Semiconductor memories: density and diversity
by Tom Williams—As CMOS technology comes into its own, system performance bottlenecks can be broken. Moreover, the hunger of modern computer systems for memory makes it practical to produce chips that meet unique needs.

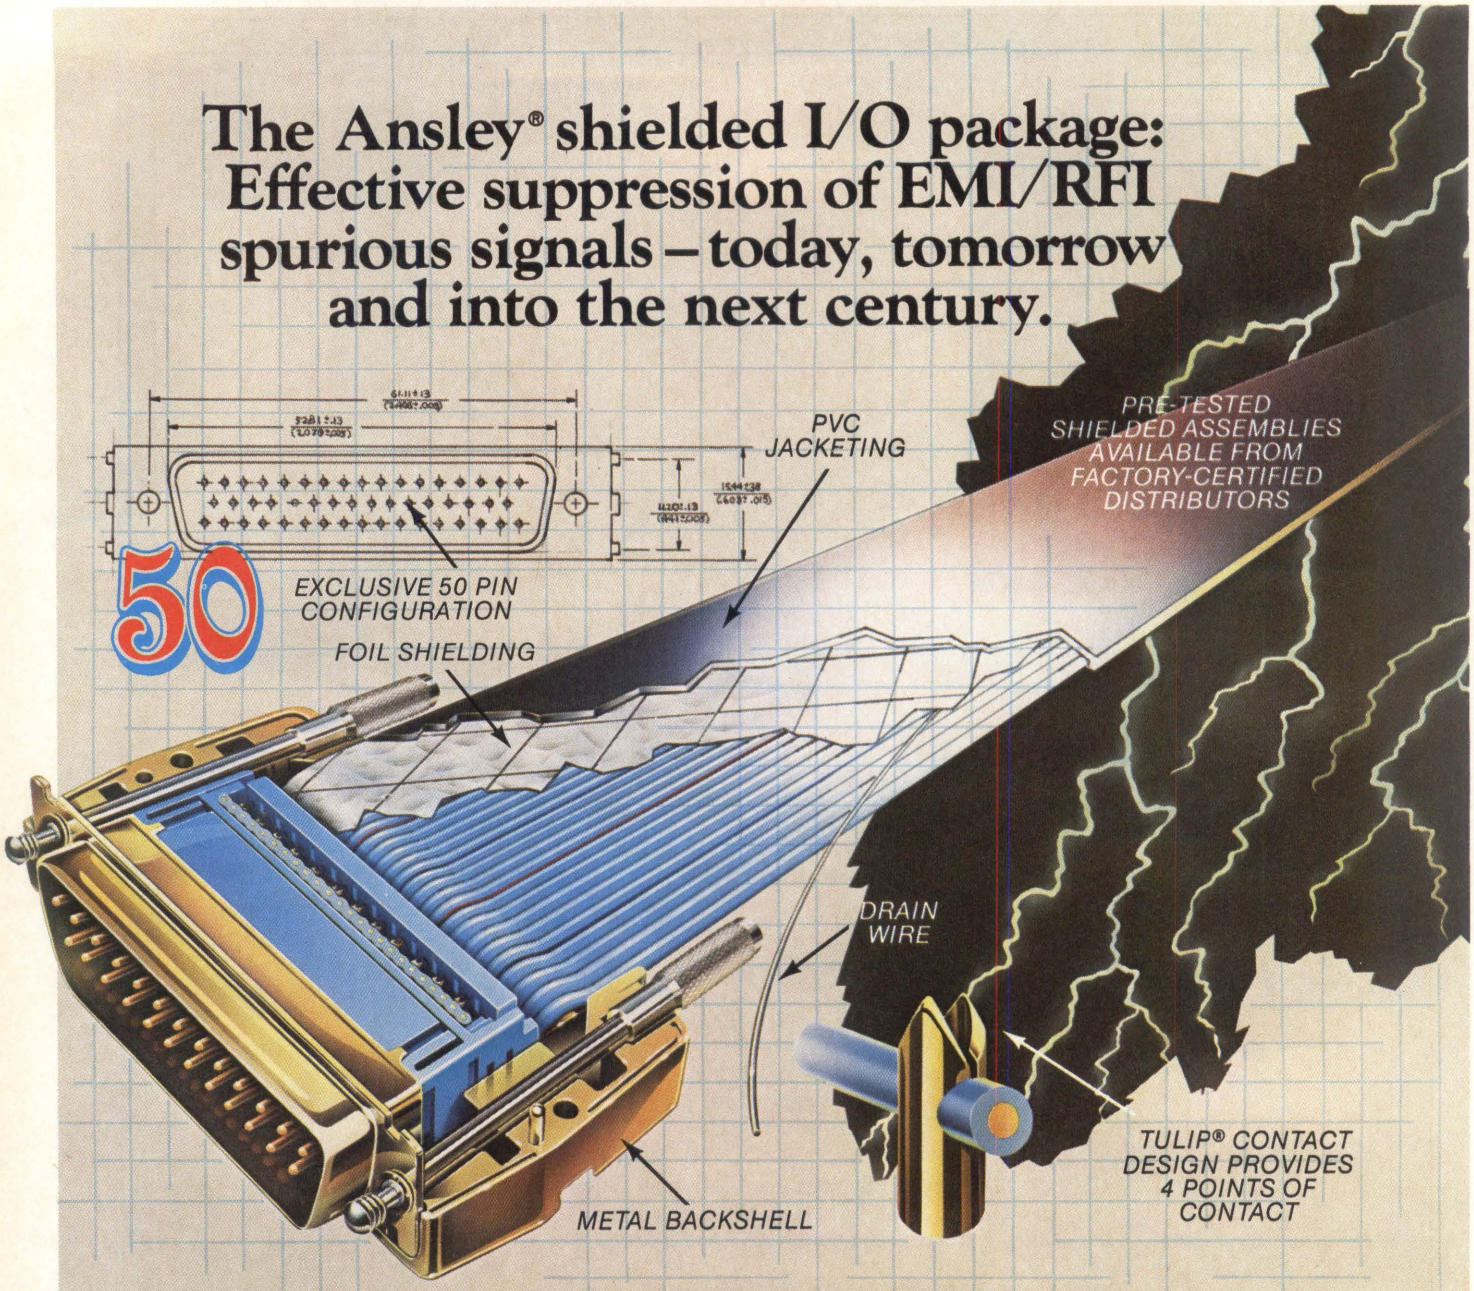
121 Joining text and graphics enhances video performance
by David W. Gulley—A dual-port RAM with a built-in shift register eliminates bottlenecks and speeds data transfers.

133 CMOS 256-Kbit RAMs are fast and use less power
by William H. Righter—By using a high performance CMOS dynamic RAM technology, relaxed timing margins combined with faster cycles can yield significant improvements in system speed.

145 Dual-port static RAMs can remedy contention problems
by Michael J. Drumm, James B. Harris, and Michel Ebertin—System bandwidth can be used up by contention problems. A true dual-port RAM can help by allowing simultaneous reads and writes by different processors.

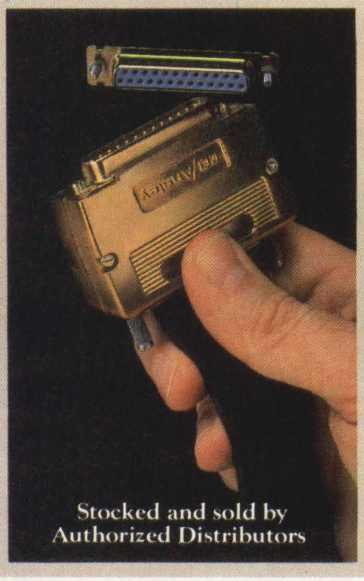
155 High endurance EEROMs open system opportunities
by Charles Furnweger—Because they can endure more read/write cycles, the latest electrically erasable ROMs can handle a wider variety of applications.

The Ansley® shielded I/O package: Effective suppression of EMI/RFI spurious signals – today, tomorrow and into the next century.



Don't let interference get in the way of a good design. Meet the only shielded flat cable and metal shell I/O connectors with Ansley reliability. They're designed to help you comply with the latest FCC EMI/RFI emissions directive, with interconnection integrity to last well into the 21st century.

Our Pin and Socket Metal Shell "D" Connectors combine reliable shielded performance with the mass-termination integrity of our patented TULIP® contact. Choose from standard 25- and 37-position types, or the industry's first mass-terminated 50-position metal-shell "D" connector. Plus a full complement of metal backshells, strain reliefs and installation hardware.



Complete your I/O interface with Ansley® Jacketed Shielded Flat Cable – famous for uniformity, electrical integrity and effective EMI/RFI suppression. And available in 9 through 64 conductors.

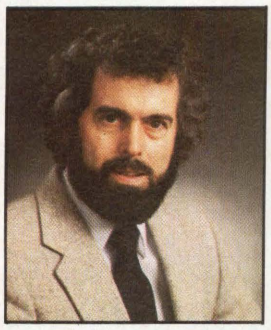
At Ansley Electronics, we're putting spurious signals in their place . . .and keeping them away from yours. For complete data, contact Customer Service, Thomas & Betts Corporation, Ansley Electronics Division, 920 Route 202, Raritan, N.J. 08869; (201) 469-4000.

T&B

Ansley Electronics Division

Where Interconnection Meets Imagination.

CIRCLE 59



SPECIAL REPORT ON

SEMICONDUCTOR MEMORIES

Advances in semiconductor memory technology are fueling an industry that, in turn, is spurring greater memory advances. CMOS breakthroughs not only offer advantages in power consumption, density, and reliability for present systems, but open a new arena for portable computer systems, as well as data acquisition equipment, and instrumentation. This will drive memory development even further as it becomes practical to tailor memory chip features to specific system needs. Far from just being parts that store more data and access it faster, memories, by their very diversity, are offering solutions to system-level problems.

Graphics systems are prime examples. The fact that they demand large amounts of memory has made it practical for manufacturers to design chips with specific graphics display needs in mind. Paramount among these needs is how to keep graphics memory updated while reading data out to the display. This section contains several approaches to the problem using specially designed memories. Even the dynamic RAMs used in computer main memories are developing more efficient addressing modes for certain situations. These include static column, nibble, and page modes that allow fast access to successive bits.

Nonvolatile storage is becoming fast enough and dense enough to graduate from its position as the repository for bootstraps, parameters, and setup tables to containing full operating systems and application software. Electrically erasable PROMs promise to make systems more flexible and to reduce the tremendous burden of updating firmware in the field. By acting as the cache memory and control store, static RAMs are becoming fast enough, and in CMOS, economical enough, to match microprocessor speed to the rest of the system. In the control storage role, they offer the possibility of more flexible microcode in microprocessor-based systems.

Lastly, memory technology is not confined to separate ICs. It is migrating onto VLSI subsystems in the form of onchip EEPROM or static cache. Because of the versatility, density, speed, and convenience of today's memories, there are endless possibilities for system innovations.

Tom Williams
West Coast Managing Editor

ROCKWELL SEMICONDUCTOR TECHNOLOGY ACCELERATES DELIVERY OF CMOS ROMS.

Rockwell's R23C64 CMOS Static ROM is right on time for your portable equipment applications.

Rockwell International is in production and accepting codes now for our R23C64 CMOS memory—one of the latest members of our CMOS product family.

This mask-programmable 64K CMOS ROM has an access time as fast as 150 ns, yet power dissipation is extremely low—10mW active, 50μW passive—so it's ideal for your low-power application requirements.

And with our 24-hour code-approval process and competitive lead times, you get high-speed CMOS without a lot of waiting.

In addition, the R23C64 is housed in a 28-pin JEDEC standard (B version), so it's pin-compatible with 64K CMOS EPROMs—allowing for easy transition with the benefit of lower power and higher speed.

For a fast solution to your low-power requirements, call your local Rockwell sales representative today. He'll tell you where to send codes and how to order low-power, high-speed CMOS memories now from—

Rockwell Semiconductor Products Division

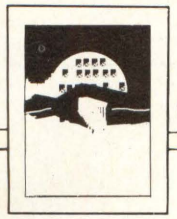
Rockwell International, P.O. Box C, MS 501-300, Newport Beach, CA 92660. Call Toll Free (800) 854-8099. In California, (800) 422-4230.



Rockwell International

...where science gets down to business

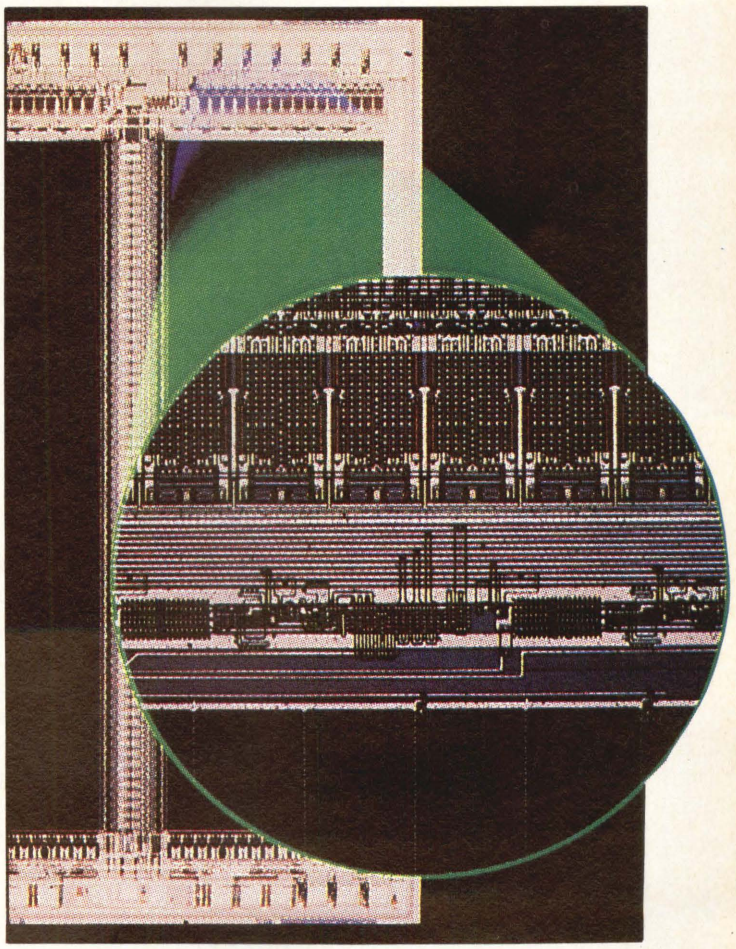
CIRCLE 60



SEMICONDUCTOR MEMORIES: DENSITY AND DIVERSITY

As CMOS technology comes into its own, system performance bottlenecks can be broken. Moreover, the hunger of modern computer systems for memory makes it practical to produce chips that meet unique needs.

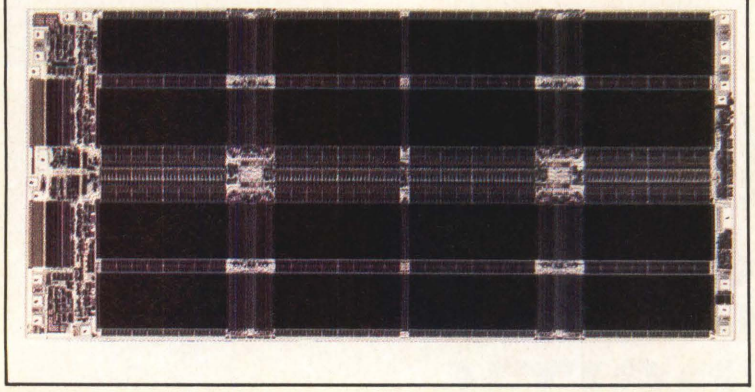
by Tom Williams,
West Coast Managing Editor



"All of our future designs will be in CMOS." This phrase echoes in the halls of semiconductor manufacturers as if it were a new litany of the electronics industry. Having overcome restraints in speed, latchup, yield, and complexity, CMOS technology is now the impetus behind the driving force of digital ICs—memory technology. No doubt, the two biggest factors affecting semiconductor memory technology today are the coming of age of CMOS, and the migration of memory devices into quasi-specialized application areas. The latter development is dependent on the circuit design type made possible by CMOS, and the sheer size of the memory market.

No memory technology discussion can afford to ignore the influence of the marketplace on the kinds of devices being developed—their functions, structure, and organization. CMOS remains the basis and the wellspring—the *sine qua non* without which today's advances in memory would not be happening.

As the market drives manufacturers toward greater densities and more complex circuits, the transition from dominant NMOS technology to CMOS is becoming a natural one. Admittedly, although the CMOS process is more complex than NMOS—it can require from 8 to 11 mask layers—circuits implemented in CMOS are simpler, requiring fewer



The Intel 51C256 dynamic RAM is representative of the new generation of 256-Kbit DRAMs. Implemented in the company's CHMOS process, the chip features Ripplemode for extremely fast access to successive bits.

transistors than their NMOS equivalents. In addition, memory circuits in CMOS can be implemented with less than half the internal clock generators of NMOS, and require significantly fewer and simpler periphery circuits.

Manufacturers are willing to accept the rigors of a more complex process in order to gain other benefits, such as circuit density and simplicity. Larger wafers (newer companies are installing 6-in. fab lines) and more reliable techniques have all but canceled the feared impact of smaller die size and device complexity on yield. In some cases, manufacturers have even used two layers of metal for interconnection. This technique has greatly reduced chip layout headaches. But, it is just the beginning.

CMOS paves the way

The advantages of CMOS extend to those very dreams that have long lived in the minds of designers. Known for its modest power consumption, CMOS has paved the way to high density chips that consume small amounts of power at low operating temperatures. CMOS speeds now rival those of bipolar circuits, while still maintaining a power consumption and temperature advantage. Especially important for memory technology, CMOS circuits can be produced that are immune to soft errors caused by alpha particles, and that preserve noise immunity as geometries are scaled even smaller.

From the system's point of view, CMOS memory technology opens a new world of portable and battery-powered systems—not only computers, but also instruments and data acquisition devices. Market size has spawned a diversity of memory device designs that make the "by one" organizations of yesteryear only one among a host of options. Dynamic and static RAMs with a wide range of power/speed combinations, and cell organizations aimed at different general application areas, are appearing. Byte-wide, nibble-wide, and "by nine" schemes have appeared, as have DRAMs with onchip refresh, and SRAMs with onchip batteries.

As a result of the huge appetite that modern systems have for main memory, today's memory design is driven by the requirements of the system in which it is to be used. This was not always the case. But, now that main memory is a commodity that can be made at a very low cost, some portion of system cost can be devoted to more specialized memories. And, although these memories are more expensive individually, they enhance system performance and contribute to its overall value. Examples of this are the control store and the cache, which are used to match processor speed to main memory speed.

A system designer must consider many factors when choosing an appropriate memory for a system: reliability, speed, cost, power consumption, size, and ease of use. But, paramount among these is reliability—not a single application area exists where reliability is not the number one priority.

Actually, American computer manufacturers are beginning to imagine perfection (eg, zero failures) in memory parts. Hewlett-Packard (Palo Alto, Calif) is well-known for working with its memory vendors to ensure that only fully functional parts are received. Ultimately, the vendor produces only fully functional parts. A number of companies have set up joint parts qualification programs with their suppliers to qualify new parts, and have established quality circle programs with vendors to ensure ongoing improvement of delivered chips. The industry's current sentiment is that this cycle can eventually result in 100 percent perfect deliveries, thus removing reliability as an issue.

A summary of the memory requirements in computer subsystems shows that among the ranked requirements, reliability is not listed because it is always number one. By far, the largest consumers of memory are main memory and graphics memory, and in both cases, the leading ranked requirement is cost. For this reason, the technology of choice is the DRAM. Performance is second because main memory speed greatly influences overall system performance.

This is also true in graphics subsystems that are becoming increasingly larger, and migrating into the personal computer arena as separate subsystems. In fact, one spokesman for Microsoft (Bellevue, Wash) says that in a couple of years any computer costing over $1000 that does not have a bit-mapped display will not be worth it. Thus, while the dollar amount is arguable, it is clear that graphics are demanded by the market, within economical reach of a large market, and thus will be large consumers of memory. It is estimated that graphics applications alone will consume fully one-quarter of total DRAM production. Such volume is currently moving manufacturers to design DRAMs specifically for graphics systems with high volume parts.

The ease of use consideration depends mainly on the type of system involved. For example, a large main memory subsystem would most likely use dense DRAMs, external refresh, and control circuitry. This makes sense both from the point of view of cost and design simplicity. On the other hand, a smaller system might be better served by using a DRAM that incorporates refresh circuitry onchip, as does the 8-K x 8 2186 iRAM by Intel (Santa Clara, Calif), or a 32-K x 9 iRAM being developed by Hitachi (Tokyo, Japan). Ease of use is also a consideration when a fairly small amount of memory is being incorporated for things such as line buffers and peripheral control ROMs. Minimizing subsystem complexity is then an important factor.

Large systems require high speed cache and control store memory. Thus, because size is not a big factor, cost need not be either. The main requirement is speed, and the fastest RAMs find their way into these applications. Previously, this meant bipolar memory, but soon, CMOS will debut in many areas, especially as higher end systems become microprocessor based.

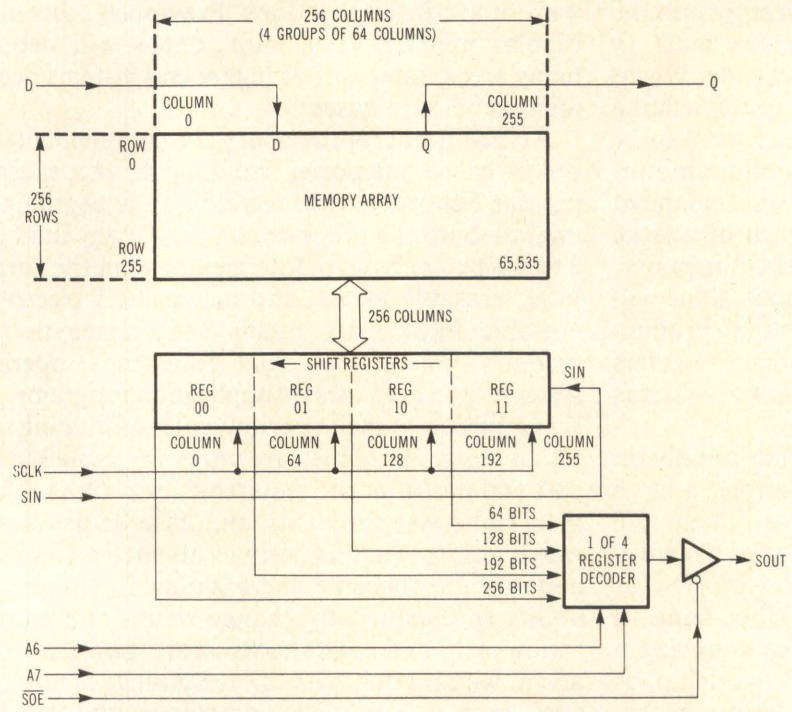
As 32-bit microprocessors are complemented by CMOS cache memories running at the speed of bipolar Schottky, that technology is beginning to migrate onto the processor chips, usually in sizes of 2 to 4 Kbytes. Nonvolatile memories in the form of ROM, erasable PROM, and increasingly electrically erasable PROM, are taking their place in main memory. There, they hold bootstraps, operating systems, and even canned application programs. This is in addition to their previous roles of holding setup parameters. Here, they are under the same kinds of cost and performance requirements as DRAMs. Used in peripherals, simplicity and ease of use are the watchwords. And, as peripherals such as terminals and printers become increasingly intelligent, the ability to dynamically change setups and configuration will make EEPROMs more popular. Once again, we find that memory technology is migrating onto such ICs as disk controllers and Ethernet controller chips.

DRAMs: density and cost

No sooner had the 64-K DRAM become economical and widely used, than the world began looking for the 256-K DRAM. But, the development of 64-K DRAMs marks a kind of watershed. For one thing, parts of the first designed in NMOS are converted to CMOS. This is not new in itself, but the

	CPU CONTROL STORE	CACHE MEMORY	MAIN MEMORY	TERMINAL/ PRINTER	GRAPHICS SYSTEM
PERSONAL COMPUTER	–	–	X	X	X
BUSINESS COMPUTER	–	–	X	X	X
SUPERMINI	X	X	X	X	X
MAINFRAME	X	X	X	X	X
					BECOMING MANDATORY
SIZE	2 to 32 KBYTES	2 to 32 KBYTES	64 KBYTES to 16 MBYTES	2 to 16 KBYTES	16 KBYTES to 16 MBYTES
RANKED REQUIREMENTS	PERFORMANCE POWER EASE OF USE COST SIZE	PERFORMANCE POWER EASE OF USE COST SIZE	COST PERFORMANCE SIZE POWER EASE OF USE	EASE OF USE COST PERFORMANCE SIZE POWER	COST PERFORMANCE SIZE POWER EASE OF USE
TECHNOLOGIES	FAST SRAM BIPOLAR PROM FAST DRAM	FAST SRAM	DRAM ROM/EPROM	MEDIUM SPEED SRAM ROM/EPROM	DRAM/VIDEO RAM

For different parts of a computer system, a different set of priorities comes into play. These needs are being met by a growing variety of memory components. Source: Texas Instruments.



As seen in the block diagram of the TMS4161, up to 256 bits of multiport memory can be shifted in parallel into the shift registers. Data can then be serially clocked out, leaving the array to be accessed randomly at any time, excluding actual parallel shifting.

diversity of system requirements and the density now possible on a single chip made it feasible to provide alternative organizations, primarily x1 and x4. The x1 arrangement was best suited for high capacity main memory needs, while the x4 made possible a reduced chip count on equipment that did not always require a full 64 Kbytes; 16 Kbytes with only two ICs.

Certainly, this trend is continuing as the industry gears up for 256-Kbit parts and beyond. The 256-K DRAMs already announced include the MCM6257 by Motorola (Austin, Tex), the MK4556 by Mostek (Carrollton, Tex), and the 51C256 by Intel—all x1 organizations—and the Mostek MK4856, which is organized as 32-K x 8 bits. National Semiconductor's (Santa Clara, Calif) 256-K x1 part, the NMC41257, uses an NMOS array with CMOS peripheral circuitry. All parts are "mainstream main memory," in that they are optimized for cost in computer main memory, which typically uses DRAMs with 120- to 150-ns access times. Interestingly enough, National has said that it does not plan to introduce a full CMOS 256-K DRAM, but that densities higher than that will make CMOS imperative. An example of the kind of scaling that has taken place already is found in the fact that a 256-K cell will fit within the width of a 16-K contact.

As for future capacities, the current horizon (we will not speak of "limits") being discussed is 4 Mbits. A number of 1-Mbit designs were discussed at the 1984 International Solid State Circuits Conference in San Francisco, and IBM (Armonk, NY) recently announced that it had produced a working 1-Mbit DRAM in the laboratory with a 150-ns access time.

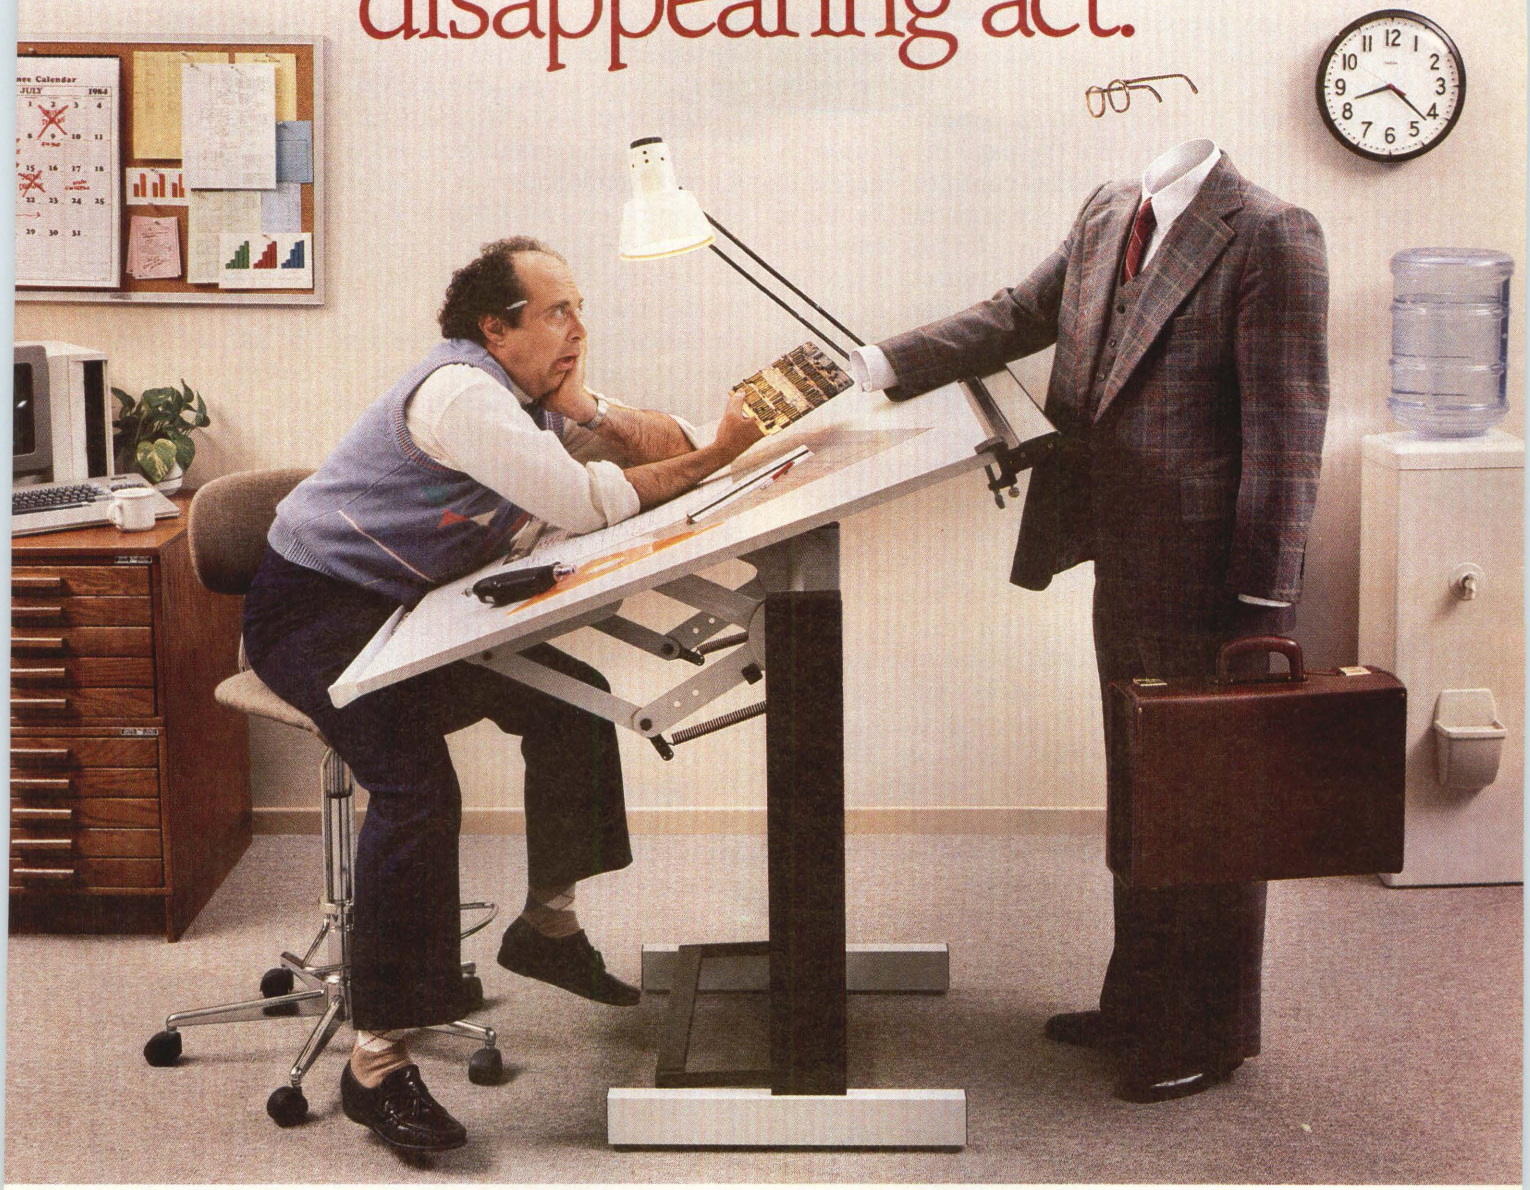
Graphics memory requirements

Another big factor, starting with the debut of 64-K DRAMs, is the desire to design chips that meet the special demands of graphics systems. Graphics is a memory intensive area and, as with main memory, cost is of utmost concern. But graphics displays need the ability to update the display memory almost as fast as it can be read. Because of this, graphics memory must deal with contention for the memory bus between the processor and the display. Among the approaches available are dual-port RAM, chips, memories with onchip shift registers, and chips offering static column and page access modes.

Some methods allow the parts to be used as any x1 DRAM. But, there are methods for reading and writing successive blocks of data at a faster rate than the normal per-bit access time. These include static column mode, page and nibble mode, and Intel's own "Ripplemode." Static column mode is a faster method of addressing successive bits. It addresses the first bit in the normal manner taking the normal access time. After this first access, the row address strobe ($\overline{RAS}$) and column address strobe ($\overline{CAS}$) signals are kept low, and the user need only increment the address signals to access the next bit. The only difference for a write operation is that the write signal must qualify the address and data in.

Nibble mode allows fast operation on 4 bits. With the first of the 4 bits addressed in the normal

We won't do a disappearing act.



Naturally every controller company talks about the quality, the reliability and the performance of its products. But few pay enough attention to an equally important issue: Responsiveness to the customer.

Fact is, just getting the kind of service you need can be a a tricky proposition these days. And, unfortunately, if you're not one of the industry's largest customers, you'll probably see that service disappear when you need it the most.

That's why we designed our Responsive 1 program. It's more than just a slogan. It's a whole new act we've built to give you the best products *and* the best service. More rapidly than ever before.

Responsive People

Headlining our new act is a cast of highly qualified marketing and technical professionals. People who'll give you the most responsive support in the industry.

In supporting roles, we've staged a worldwide product distribution network with reps and distributors who get involved with what you're doing. And who know what they're doing.

Responsive Products

For an encore, we've spotlighted engineering excellence. The kind that's introduced almost two dozen new controller products over the last five years.

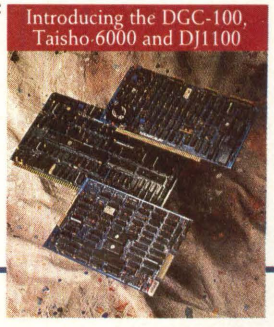
We've engineered feature packed Multibus* and S-100 products, and designed our own David Junior general purpose controllers. Now we're debuting an ensemble of star performers featuring multiple form factors, floppy and QIC backups, SASI/SCSI bus standards, custom VLSI circuitry, and much more.

So before some controller company does its disappearing act for you, contact the company that'll give you a solid performance. Konan Corporation.

The responsive controller company.

KONAN™

The Responsive 1.



1425 N. 27th Lane/Phoenix, AZ 85009/(602) 269-2649/TELEX 910-951-1552 *Trademark of Intel Corporation

CIRCLE 61

way, $\overline{RAS}$ can be kept low, and $\overline{CAS}$ cycled up and down to read the next 3 bits. Row and column addresses are supplied on the first access of the cycle. After that, the falling edge of $\overline{CAS}$ activates the next bit. After accessing 4 bits, nibble mode wraps around to the first bit accessed.

Intel's Ripplemode, like the page mode used by other manufacturers, provides continous access to successive bits. Ripplemode accesses the first bit in the normal manner. After that, $\overline{RAS}$ is kept low, and the $\overline{CAS}$ line cycled. But, the address is latched on the leading edge of $\overline{CAS}$ and the data is output on the falling edge. This provides a kind of "look-ahead" or pipelining of the addressing. Page mode is similar, but does not use Ripplemode's lookahead technique. Instead, the data out responds after the address signal to $\overline{CAS}$ has propagated through the chip. Thus, Ripplemode is somewhat faster.

Another way of serving the needs of graphics systems is to use a multiport RAM as in the case of Texas Instruments' (Dallas, Tex) TMS4161. The TMS4161 incorporates one port for random access. It also contains four cascaded 64-bit internal shift registers that output at a 25-MHz shift rate via a fast serial port. With this, the registers can be loaded after every 64-, 128-, 192-, or 256-shift cycle, depending on how the user configures them. Data can be written into the RAM via the random access port, at the same time it is being shifted out the serial port. The only time the array is unavailable for updating is during the loading of the shift register(s). Articles that include examples of both approaches to graphics subsystem design appear in this feature section.

High speed SRAMs

From the system designer's point of view, SRAMs, especially the high speed variety, are becoming the memory of choice where performance is the prime requirement. Since SRAMs require more transistors per bit, they cannot be made as dense or at prices per bit comparable to DRAMs. Therefore, they are less likely to be used in large main memories. Their preferred application is in small, portable systems, and in larger systems for CPU control store and cache. In multi-user systems, the cache memory represents the performance bottleneck, as does the control store memory, in terms of the CPU's ability to reach its optimal performance. Here, speed is most important, and the amount of memory required is small compared with that of the whole system.

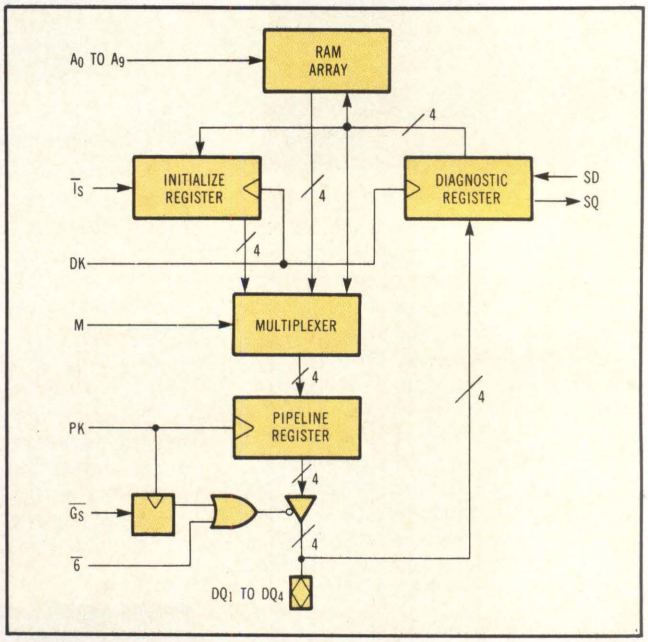
Although bipolar ROM has often been the memory of choice for control store, the availability of fast SRAM is creating enticing possibilities. Ramtek (Sunnyvale, Calif) recently introduced a new graphics system, the microprocessor-based 2020, which uses fast SRAM for writable control store. The company supplies the microcode on a 3½-in. floppy diskette. The ability of fast CMOS SRAM to match the speed of the microprocessor has allowed the system to use an older method of storing microcode, and provide a new level of flexibility in the system.

Fast CMOS SRAMs seem to be part of an area where smaller manufacturers are making inroads. Lattice Semiconductor (Portland, Ore) is expected to announce a 64-K CMOS part with a 35-ns access time, and Cypress Semiconductor (San Jose, Calif) is pushing 12-ns speeds with a 256 x 4-bit CMOS SRAM, the CY2122.

Basically, there are three kinds of SRAMs: commodity parts, for such things as portable systems where they run in the 100-ns and slower range; high performance parts, pushing the speeds mentioned above; and certain specialized SRAMs. One example of the latter is the Am9151, a 1-K x 4 CMOS SRAM with a 40-ns cycle time from Advanced Micro Devices (Sunnyvale, Calif).

The Am9151 incorporates three high speed 4-bit registers: a parallel pipeline register; a shiftable shadow register; and an initialize register used to generate any arbitrary microinstruction for system interrupt or reset. The shadow register controls and observes the pipeline register during a diagnostic or test mode. It can be used in loading a writable control store by serially shifting an instruction word into the shadow register and then clocking the data in parallel into memory.

In effect, previously inaccessible internal state information can be shifted out via the serial data output and observed. These serial data paths can



The Advanced Micro Devices Am9151 static RAM contains three internal registers that enable the designer to monitor internal states of combinational logic blocks for diagnostic operations.

CMOS is a Solid State of mind.



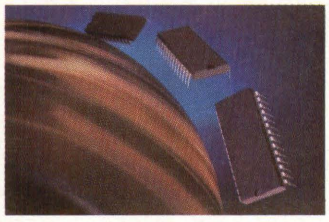
New 128K CMOS ROMs from Solid State Scientific.

Fast access times (100 ns max.)
Faster system execution.
Fastest turnaround times in the business.

Solid State Scientific has a new 128K CMOS ROM. The 23C128. It's extremely fast—75 ns typical or 100 ns worst case. So you can use it with the world's fastest microprocessors. It has the low power and high reliability of our state-of-the-art CMOS technology. It's available now at prices that make it as cost-efficient as it is technically superior. Production volumes take only 7 weeks, prototypes just 3 weeks.

Forget your NMOS ROMs and EPROMs

If you're using NMOS ROMs, our new 23C128 can give you far better speed and reliability. If you're using EPROMs, our pre-programmed fast ROMs can speed your system execution by eliminating wait states. With much faster turnaround, too. And, at a cost more than 25% below EPROMs.



Quick specs on the 23C128

3 Speeds: 150, 120 and 100 ns max. **Standby Current:** 50 microamps max. **Operating Current:** 10 milliamps max. LSTTL-compatible inputs and outputs. 28-pin JEDEC standard. Military versions, too.

Much more CMOS memory coming from SSS

In case you didn't know it, SSS delivered more 32K and 64K CMOS ROMs than any domestic supplier last year. And we're out to do the same with our new 128K CMOS ROM. Other new CMOS memory devices, including a new 256K CMOS ROM, are on the way.

So why wait?

For details on our new 23C128 CMOS ROM or any of our other CMOS memories or logic devices, just call or write. And find out why CMOS *is* a Solid State of mind. Solid State Scientific, 3900 Welsh Rd., Willow Grove, PA 19090. (215) 657-8400.

Solid State Scientific
CMOS ResponsAbility™

CIRCLE 62

also be cascaded. In addition to control operations, this serial shadow register diagnostic technique can be used with such important state registers as macroinstruction, status, data, and address. System diagnostics become easier by breaking the normal feedback and turning sequential state machines into combinational logic blocks whose internal states can be monitored.

The range of speeds and densities available in mask-programmable ROMs and EPROMs has taken them from the lowly status of holding setup and configuration data for peripheral or bootstrap programs, to the ranks of the carriers of entire operating systems, realtime kernels, and canned application programs. Part of the decision whether or not to use a ROM or an EPROM will always be dictated by cost. ROMs in volume will continue to have a price edge over EPROMs.

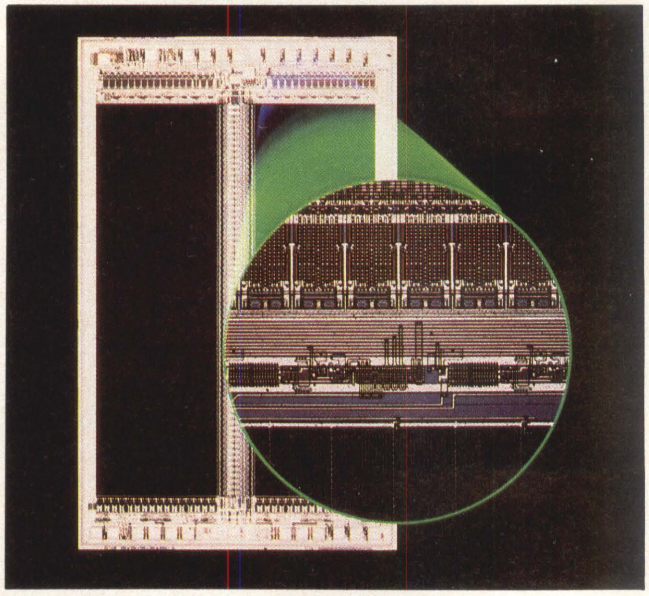
But beyond that, EPROMs are comparably priced and offer the advantage of being reprogrammable. They are often chosen for similar tasks, especially when volume places their cost below that necessary for a mask-programmable ROM. An example is Am27512, a 64-K x 8 device announced by AMD.

For those peripheral applications requiring relatively small numbers of setup parameters to be stored and changed only occasionally, specialized parts are coming into use, such as the Mostek MK48C02 2-K x 8 battery backup RAM. This part contains onchip lithium cells and voltage sensing circuits that switch on the batteries when external power is removed. Likewise, the NOVRAMs or "shadow" RAMs pioneered by Xicor (Milpitas, Calif) combine a normal RAM array with a corresponding EEPROM array. When the chip is powered up, the contents of the EEPROM are written into the RAM, providing a default configuration. Parameters can be changed while the system is on, and the "shadow" reprogrammed to change the default setting. The leading edge in NOVRAM is also pushing beyond DIP switch replacement with a 4-K (512 x 8), 300-ns part expected to appear soon.

Toward EEPROM standards

Major issues affecting EEPROMs as they move to the higher densities include providing a means of writing the memory in a reasonable amount of time; protecting the device against false writes during power-up and power-down; and considerations on the definition of pin 1 in the 28-pin Joint Electron Device Engineering Council (JEDEC) standard package. Read access times are also quite respectable, such as in the 64-K CMOS EEPROM by Exel (San Jose, Calif), which boasts 55 ns.

The pin 1 issue appears to be the most easily resolved. Some early 64-K announcements define pin 1 as a ready/busy signal to send an interrupt to the microprocessor when the write cycle is complete. However, this may change. While ready/busy was popular in 16-K versions using a 24-pin package, the 64-K designs are looking to preserve their 28-pin site for anticipated 256-K versions. Hence, they have discovered a different method of signaling the microprocessor and reserving pin 1 for an address pin (A14) for 256-K designs.



A full 64 Kbytes can be stored in the Am27512 512-Kbit EPROM and accessed at 150 ns.

$\overline{\text{DATA}}$ polling, as introduced by Xicor, is a software method of determining if a write cycle is complete. A polling method of some sort has become necessary in EEPROMs as density increases. Given the 10-ms typical write time per byte, microprocessors would be idling too long unless they could go off and perform other tasks while the part was writing. $\overline{\text{DATA}}$ polling allows the processor to look at the data being written to see if a write is still in progress. This has two advantages. First, it frees the processor to perform other tasks during the write cycle, and second, it allows the system to optimize the actual time consumed by write operations. Since systems may incorporate large numbers of dense EEPROMs, the savings may be considerable.

It should be noted that some manufacturers (eg, Intel) have assigned pin 1 as a ready/busy pin in their 64-K, 28-pin designs. This was probably done because it anticipates designs in the 256-Kbit range to use 16-bit word widths. However, this would require a different package altogether, possibly a 40-pin type.

Increased densities of EEPROMs have made their use conceivable in an ever-expanding range of applications. The time needed to program a single byte has become a major hindrance. At 10 ms per byte, the time to program an entire 64-K part exceeds 81 s. Where smaller devices latched single bytes to free the processor, the newer ones buffer and latch groups

GOOD & PLENTY



The Fujitsu MB81416 16K X 4 DRAM. Now available in quantity.

Nobody puts tastier technology into a DRAM than Fujitsu.

Like the fastest access time available – 100ns. A TCAC/TRAC ratio of 2:1 for easier timing. Low power consumption, too.

And in a variety of packages. Like plastic DIP. JEDEC-approved compatible with the TMS 4416 and INMOS 2620.

Mmmm mmmm, good.

Best of all, we've got plenty in stock. So call the nearest FMI sales office listed below. For immediate gratification.

	100ns	120ns	150ns	Page Mode	TRAC/ TCAC	CAS before RAS refresh
FUJITSU	yes	yes	yes	yes	2:1	yes
TI	no	no	yes	yes	1.66:1	no
INMOS	yes	yes	yes	no	1.66:1	yes

FUJITSU MICROELECTRONICS, INC.

Technology that works.

3320 Scott Boulevard, Santa Clara, CA 95051 · 408/727-1700

FMI Sales Offices. Atlanta 404/449-8539. Austin 512/343-0320. Boston 617/964-7080. Chicago 312/885-1500. Dallas 214/669-1616. Houston 713/784-7111. Minneapolis 612/454-0323. New York 516/361-6565. No. Calif. 408/866-5600. S. Calif. 714/720-9688.

© 1984 Fujitsu Microelectronics, Inc.

CIRCLE 63

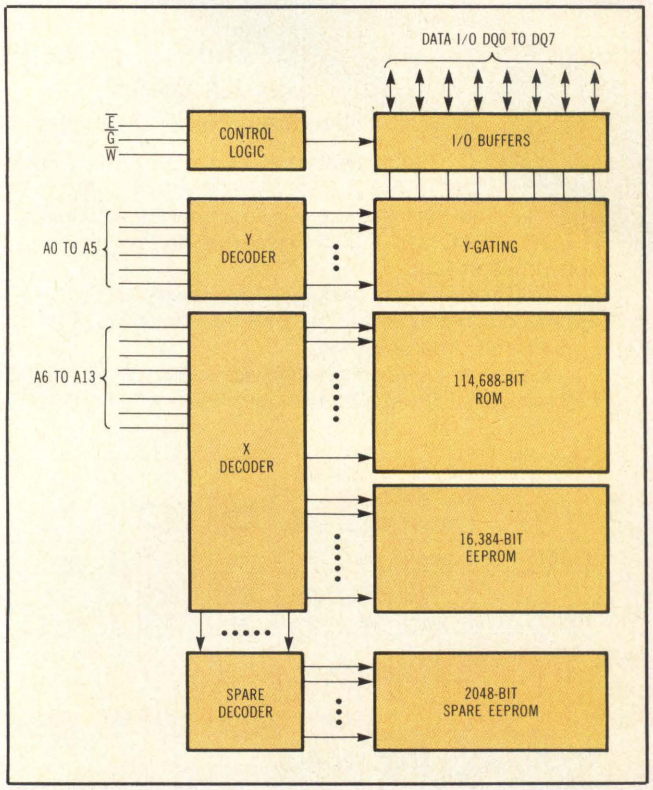
of bytes, and write each group *en masse*, while the processor is busy elsewhere. This is called page mode.

There is a trade-off to consider here also. Most applications do not require the entire EEPROM to be written over very often. Each time a group of bytes is changed to alter 1 byte, the erase/write endurance for the whole group is diminished. Although endurances are typically 10,000 erase/write cycles per byte, not all designers are entirely comfortable with that figure—especially in the new application areas developing for larger capacities. Page sizes for Intel's projected 64-K device and Xicor's 2864A are 16 bytes, while the Inmos (Colorado Springs, Colo) 48C64, and the AMD 2864 will have 32-byte buffers. A 64-K part from Inmos, the 3630, sports a 64-byte page size.

Time savings for write operations using page mode can be dramatic. Xicor, for example, specifies a 5-ms write cycle time. By latching and writing data 16 bytes at a time, the time to completely write 64 Kbits is cut to 2.5 s. Of course, EEPROMs with 32-byte pages would be half that figure.

EEPROM protection

One manufacturer, Exel, has opted to dedicate its pin 1 to a special status word (SW) function. This effectively locks its 28-pin package size out of the 256-K arena, but gives it extra software features to support page mode in its 64-K XL48C64. A set of on-chip registers can be read or written by the microprocessor when SW pin 1 is low. Thus, the chip can supply ready/busy status information to the processor when it reads a register during SW active. The processor can also write commands into the XL48C64 to set it for page mode, fast write mode, or for chip erase mode.



In addition to 16 Kbits of EEPROM in the memory map, a spare 2-K can be mapped into the memory space of Motorola's MCM683616 combination ROM/EEPROM by the decode circuitry.

Protecting an EEPROM against inadvertent writes during power-up and power-down employs a combination of voltage threshold sensors, noise filters, and time-outs. In addition to disabling write functions until memory reaches V_{CC}, it is necessary to keep the write disabled for a certain amount of time thereafter, to allow V_{CC} to settle. Intel and AMD, for example, disallow writes if V_{CC} is below 4 and 3 V, respectively. Xicor has implemented a precision V_{CC} sensor with which the user can set the threshold voltage by writing and locking an EE cell.

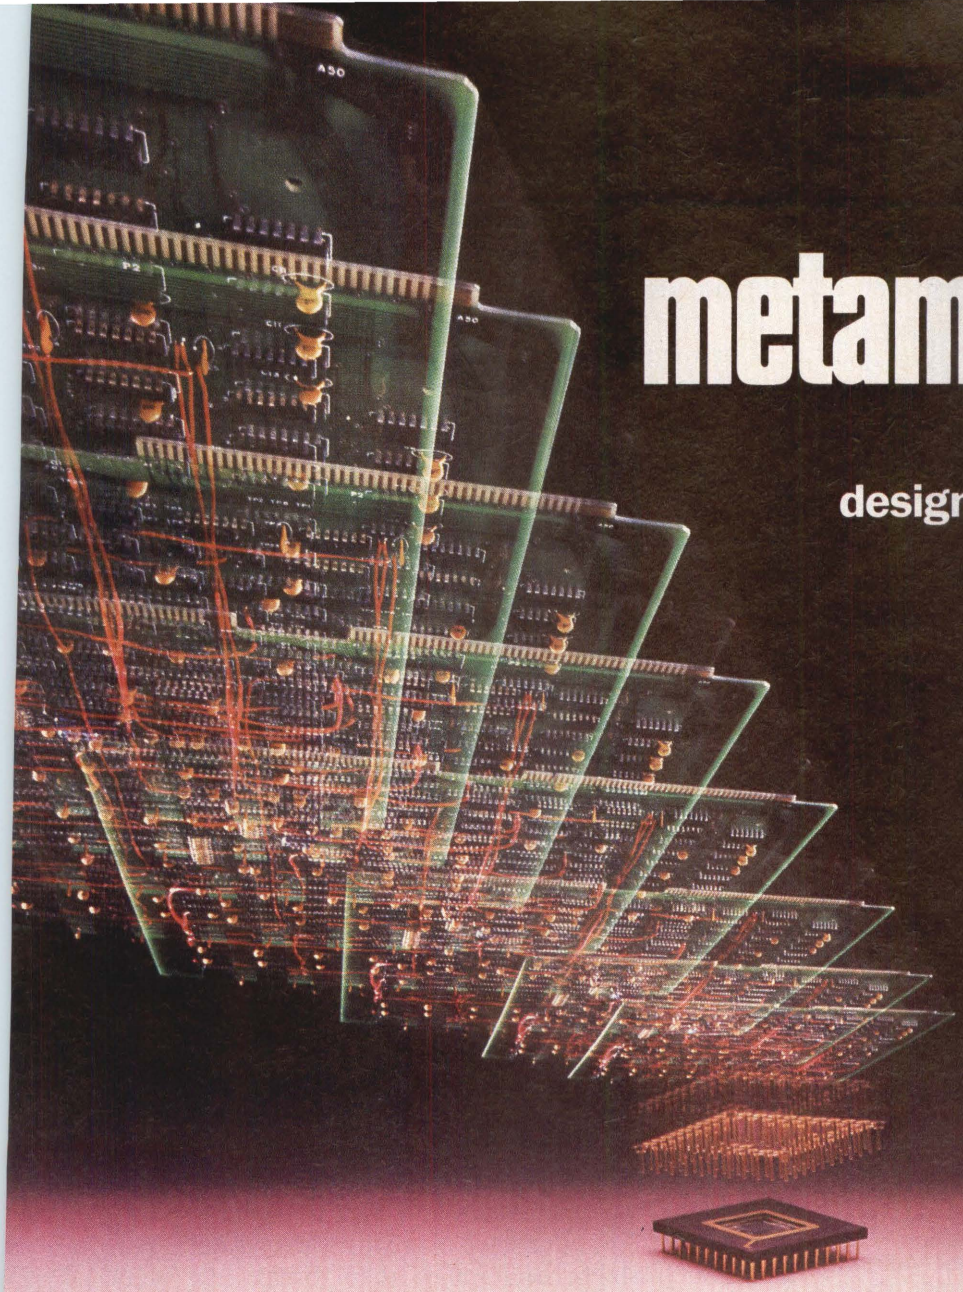
These companies also include a time-out, before the chip can be written to, of typically 100 ms on initial power-up, so that V_{CC} can become stable, and noise does not initiate an unwanted write. Exel makes use of its status word pin by requiring that a pattern be written into a register that enables the charge pump to provide the 21-V programming voltage from the 5-V supply before a write can take place. Since the registers are static, an inadvertent write on power-up cannot take place, but the chip must be initialized.

Electrically erasable, or E^2 technology, is finding its way into more than just straight memory products. E^2 arrays are used onchip to configure the 8001 Ethernet controller from Seeq Technology (San Jose, Calif), and to alter parts of the microcode on the 72720 microcomputer developed by Seeq and TI.

Another example of how E^2 technology is combining with other circuitry to solve cost/function problems is the 33128 ROM from TriStar Semiconductor (Santa Clara, Calif), and the Motorola MCM683616 combination ROM/EEPROM memory unit (CREEM). Both are described as mask-programmable ROMs with E^2 patch. They address the need to make changes in a firmware program or data without having to replace an entire ROM, or without going to the expense of using a full-blown EEPROM in a design. Thus, they are an interim solution until such time as E^2 technology is cost-competitive with mask ROMs.

Motorola's MCM683616 is a 12-Kbit combination ROM/EEPROM device. Organized as 16 Kbytes, its lowest order 2 Kbytes are bulk-erasable EEPROM, and the remaining 14 Kbytes are mask-programmed ROM. Another 256 bytes of spare EEPROM can be mapped into the address space to replace any 256-byte page of mask ROM or EEPROM on the chip.

Similarly, the TriStar 33128 is a 16-Kbyte mask-programmable ROM on which certain portions are implemented as EEPROM. This is done in two ways.



AMCC metamorphosis

rapidly reduces your designs to high performance ECL/TTL logic arrays

At AMCC, metamorphosis means a lot more than just reducing board real estate, material cost and assembly time. AMCC metamorphosis transforms your designs from net lists to prototypes in as little as four weeks. It means a dramatic semicustom evolution using AMCC's system oriented logic arrays.

AMCC metamorphosis pushes aside past limitations in high performance system applications by offering high density logic at sub-nanosecond speeds; ECL, TTL or mixed I/O on the same chip; up to 95% gate utilization; advanced MSI/LSI macro libraries; automatic place and route software; the highest commercial and military quality and reliability standards (Mil. Std. 883); Class 10 fabrication; and mask compatible alternate sourcing.

Our engineer-to-engineer full service commitment plus complete CAD capability helps provide you with fast flexible design solutions. To us full service means support . . . a design implementation group; Daisy™ and Mentor™ engineering workstations; Tegas™ via Cybernet™; field applications engineering; on-site training courses; complete documentation; portable design centers; complete wafer fabrication; and quick turn-around assembly and test.

AMCC ECL/TTL LOGIC ARRAYS

	Q3500S	QH1500A	Q1500A	Q700	Q710	Q720
Equivalent Gates	3500	1700	1500	1000	500	250
Typ. Gate Delay (ns)	.3 – .7*	.9	.9	.9	.9	.9
Typ. Power (W)	3.5	2.8	2.5	2.0	1.2	.6
I/Os	120	120	84	76	56	36
Gate Utilization	95%	95%	95%	85%	85%	85%

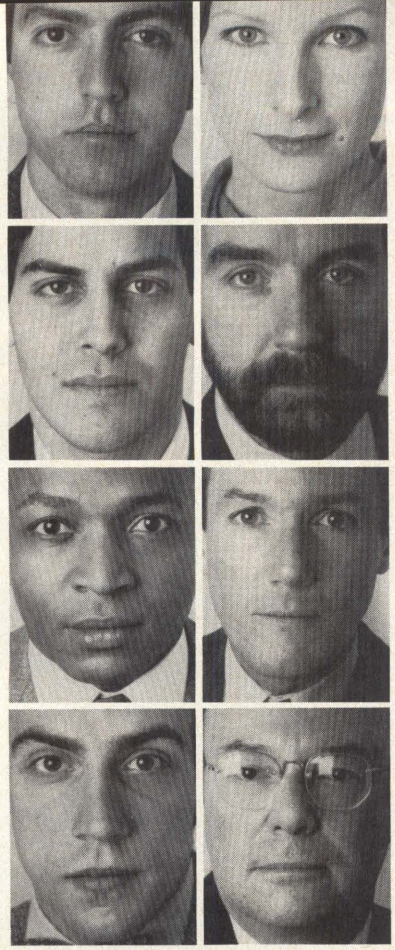
*programmable

AMCC . . . dedicated to high performance semicustom. Make us perform for you.

Call or write for full information on how AMCC's metamorphosis can provide high speed silicon solutions to your system problems.

AMCC
APPLIED MICRO CIRCUITS CORPORATION
5502 Oberlin Drive
San Diego, California 92121
Telephone (619) 450-9333 • TWX/Telex: 910-337-1136

CIRCLE 64



WHERE TO MEET THE MOST OEM'S AND SYSTEMS INTEGRATORS FACE TO FACE TO FACE...

Meet them at a Mini/Micro.

It's the one computer trade event that focuses entirely on computer design needs. That means you, as an exhibitor, can deal face to face with a prime concentration of OEM's, designers, systems integrators and software developers. On your own turf, but in their backyard; Mini/Micros are held regularly in the prime markets.

And many run concurrently with other high-tech electronics shows. That gives you extra value in the form of crossover attendees.

So if you're after OEM's and systems integrators, be at a Mini/Micro. For complete information, call toll-free: 800-421-6816. In California, 800-262-4208.

Meet OEM's and systems integrators where they live.

Sept. 11-13 1984	Mini/Micro Southwest Midcon	Dallas Dallas
Oct. 2-4 1984	Mini/Micro Northwest Northcon	Seattle Seattle
Feb. 5-7 1985	Mini/Micro West	Anaheim
Mar. 5-7 1985	Mini/Micro Southeast Southcon	Atlanta Atlanta
May 23-25 1985	Mini/Micro Northeast Electro	New York New York

Sponsored by regional chapters of
IEEE and the Electronic Representatives Association

Mini/Micro
FOCUSING ON THE OEM

116 COMPUTER DESIGN/August 1984 CIRCLE 65

First, the last 256 bytes of the address space are reprogrammable. This has a certain amount of usefulness, but is limiting in that these bytes have a fixed address. Programmers can easily change data in those addresses, but they cannot change data or instructions in the rest of the program.

There are 352 additional E^2 bytes onchip that can be mapped to any address space in the ROM area. These 352 bytes are broken into 2-, 4-, and 8-byte modules; a programmer selecting a module must map all bytes in that module. The 33128 has onchip logic to perform the mapping and can be reprogrammed in-circuit. The ability to change certain parts of a firmware program is useful for machine control applications where wear may require certain values to be updated, or for calibrating intelligent instruments in the field.

Packaging issues heat up

If cost considerations are a driving force in memory systems, their effect is being dramatically felt in component packaging. This proves the well-known truism that component size drives system cost. For 64-K devices and beyond, there is an industry trend toward the leaded chip carrier (LCC). The LCC is comparable in cost to the DIP, yet offers high density placement, 50-mil lead spacing, and easy shipping and handling.

LCCs also lend themselves to surface-mount manufacturing techniques. As surface-mount equipment is phased into manufacturing facilities, cost savings are likely to increase dramatically. The equipment is simpler than through-hole methods using wave solder equipment; surface mounting uses solder phase screeners and pick-and-place robots. In addition, LCCs have overcome the thermal expansion problems that made them break free of PC boards; their expansion characteristics have been matched to those of PC boards.

One further development in component placement made possible by the LCC is the single inline package (SIP). A SIP is a PC substrate on which four to nine LCCs and capacitors have been surface-mounted. The leads on one edge of the SIP can then be through-hole mounted into a conventional PC board. Thus, the advantages of LCCs and SIPs can be made available to customers who have not yet invested in surface-mount manufacturing equipment. A strip less than 5 in. long can hold 64 Kbytes of RAM, while using a minimum of real estate on a conventional PC board itself.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 707 *Average 708* *Low 709*

At last! Powerful solutions for identifying and eliminating datacomm and network problems.



With HP's family of protocol analyzers, you can minimize network downtime and handle even the thorniest of datacomm problems. Choose from the standard HP 4955A, or our newest member, the powerful HP 4951A, weighing in at only 14 pounds. They set new standards of performance for R&D and field service personnel involved in solving datacommunications problems, and end-users responsible for network maintenance and planning.

The HP 4951A/4955A team—your key for high-powered datacomm testing.

Here's high analytic power and productivity. The HP 4955A and HP 4951A are fully program-compatible. You can develop test programs on an HP 4955A and then transfer them to your field service personnel via a modem or tape cartridge. With the HP 4951A/4955A protocol analyzer team, you can now go wherever your datacommunications problems take you and have the confidence of on-site versatility, with back-up power just a phone call away.

HP 4955A—increase your capability with top performance features.

With the HP 4955A you can bring your products to market faster, and with a greater level of reliability than ever before. Our datacomm-enhanced BASIC lets you program sophisticated test routines and perform high-level protocol analysis. Exercise your hardware and software as it's being developed, not after it's installed. Monitor, simulate, and trigger from 50 bps to 72 kbps. You can easily identify protocol problems at the physical interface, frame, and packet levels using the HP 4955A's multiple display formats. Plus, our intelligent 256K byte buffer memory increases real data storage by eliminating line idles without sacrificing timing information.

HP 4951A—"instant" productivity.

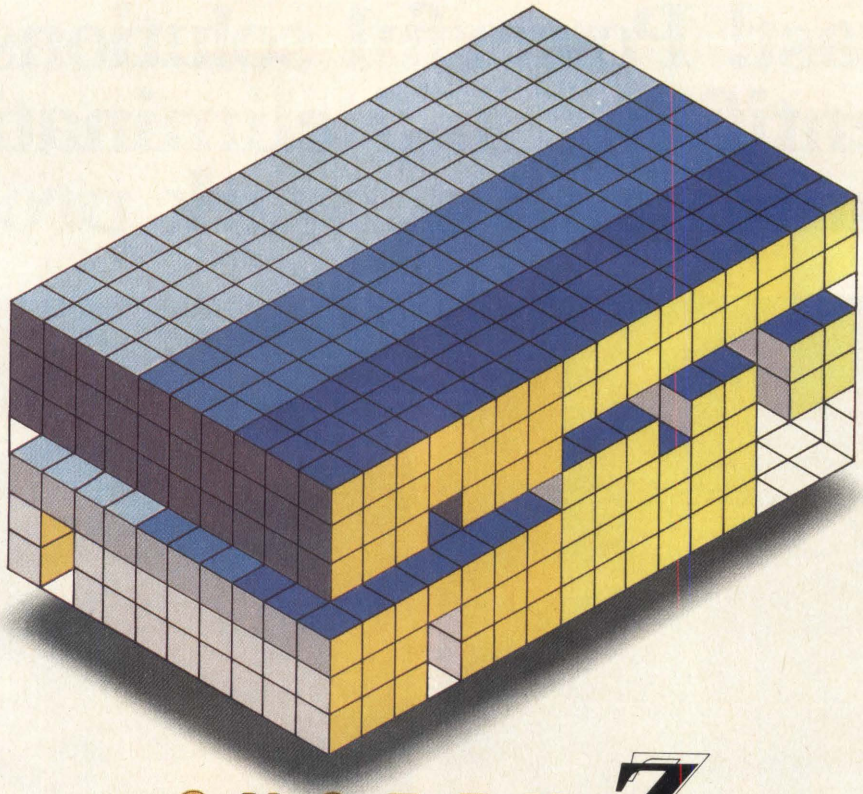
The HP 4951A features advanced one-button autoconfiguration which easily does the setup work for you and gets you monitoring data quickly on-line. Like the HP 4955A, it gives you 63 simultaneously active triggers for extensive testing. In post-processing mode, you can do detailed repetitive analysis for hard-to-track errors. You can trap on characters, error conditions and lead transitions. To isolate problems down to the network component level, BERT mode lets you measure bit errors, block errors, errored seconds, and percent error-free seconds. You can simulate a CPU, modem, terminal, or group of terminals for complete interactive testing. The HP 4951A accommodates most popular protocols, data codes, and speeds (to 19.2 kbps).

For more information on the HP 4951A and HP 4955A protocol analyzers, call your local HP sales office listed in the telephone directory white pages. Ask for the electronic instruments department.

HEWLETT PACKARD

3801401

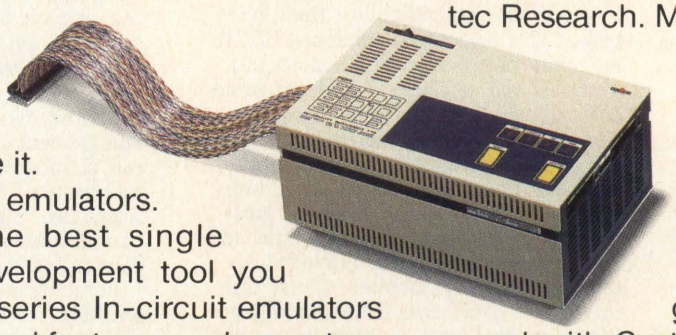
CIRCLE 66



SYSTEM Z

HOW YOU CAN OWN THE BEST MICROPROCESSOR DEVELOPMENT SYSTEM IN THE WORLD

First look to ZAX. Because ZAX knows what it takes to be the "Best." And that means offering development products that are intelligent, compatible, fast, powerful, and, what most manufacturers forget; *obtainable* – because if you can't afford to buy it, you can't use it.

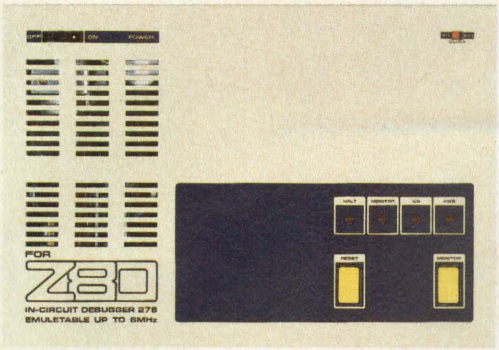


Start with ZAX emulators. They're simply the best single microprocessor development tool you can buy. ZAX ICD series In-circuit emulators offer several advanced features such as extensive emulation memory, breakpoint trace capability, true realtime emulation with no wait states, and a variety of powerful debugger commands. And as you can see, ZAX supports a broad spectrum of 8 and 16-bit microprocessors as well.

But don't stop there. To get things running, add the capable and economical IBM PC or equivalent personal computer. Now make it perform with professional software from Microtec Research. Microtec's cross-software products include symbolic debuggers, fully manufacturer compatible macro assemblers, overlay linking loaders and hi-level language support.

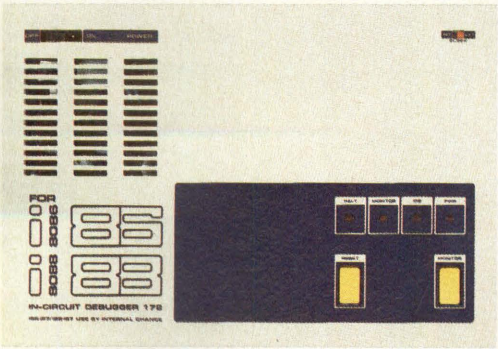
You could end there, but ZAX doesn't. As your needs grow, continued support is assured with System Z; an incredible complete development program that can link your PC with your mainframe computer.

If owning the best microprocessor development system in the world interests you, contact ZAX today at **800/421-0982** for more details.



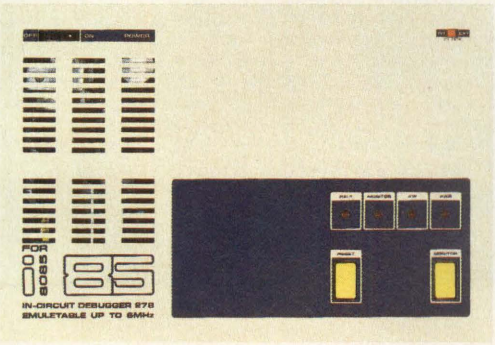
Z80
Z80
Z80B
Z80H

Emulates Z80B microprocessors to 6 MHz and Z80H to 8 MHz. Features; 64K byte user emulation memory, 2K deep x 32 bits wide realtime trace buffer, 29 different debugger commands.



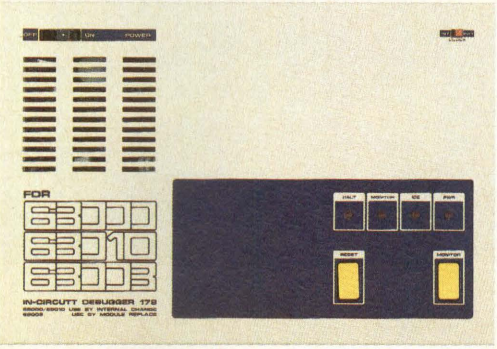
i8086/88
i8086
i8087
i8088

Co-emulation of 8086 and 8087 or 8088 and 8087 processors to 5 MHz. Realtime emulation to 8 MHz for 8086/88 processors. Features; 128K bytes static RAM – expandable to 1 Mbyte, 4K deep x 40 bits deep realtime trace buffer, 30 different debugger commands.



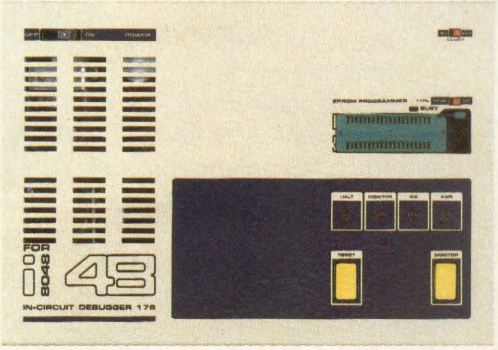
i8085
i8085A
i8085A-1
i8085A-2

Emulates 8085 processors up to 6 MHz. Features; 64K byte user emulation memory, 2K deep x 32 bits wide realtime trace buffer, 29 different debugger commands.



68000
68000
68010
68008

Emulates 68000, 68008 and 68010 in one unit to 10 MHz. Features; 128K of emulation memory – expandable to 256K, 4K deep x 48 bits wide realtime trace buffer, 30 different debugger commands.



i8048
i8048
i8049
i8050
i8748
i8749
i8039
i8035
i8040

Emulates entire 8048 family in one unit to 11 MHz. Features; 4K emulation memory, 2K deep x 32 bits wide realtime trace buffer, 29 different debugger commands. 8748 and 8749 units feature a built-in EPROM programmer.

All emulators come equipped with these standard features; 3 hardware and 8 software breakpoints, an Event Trigger and External Probe triggered breakpoint, 14 selectable baud rates to 19200bps, 20" processor probe, 115/230 VAC operation. All Zax emulators are compact, lightweight, rugged and can easily fit into a briefcase.



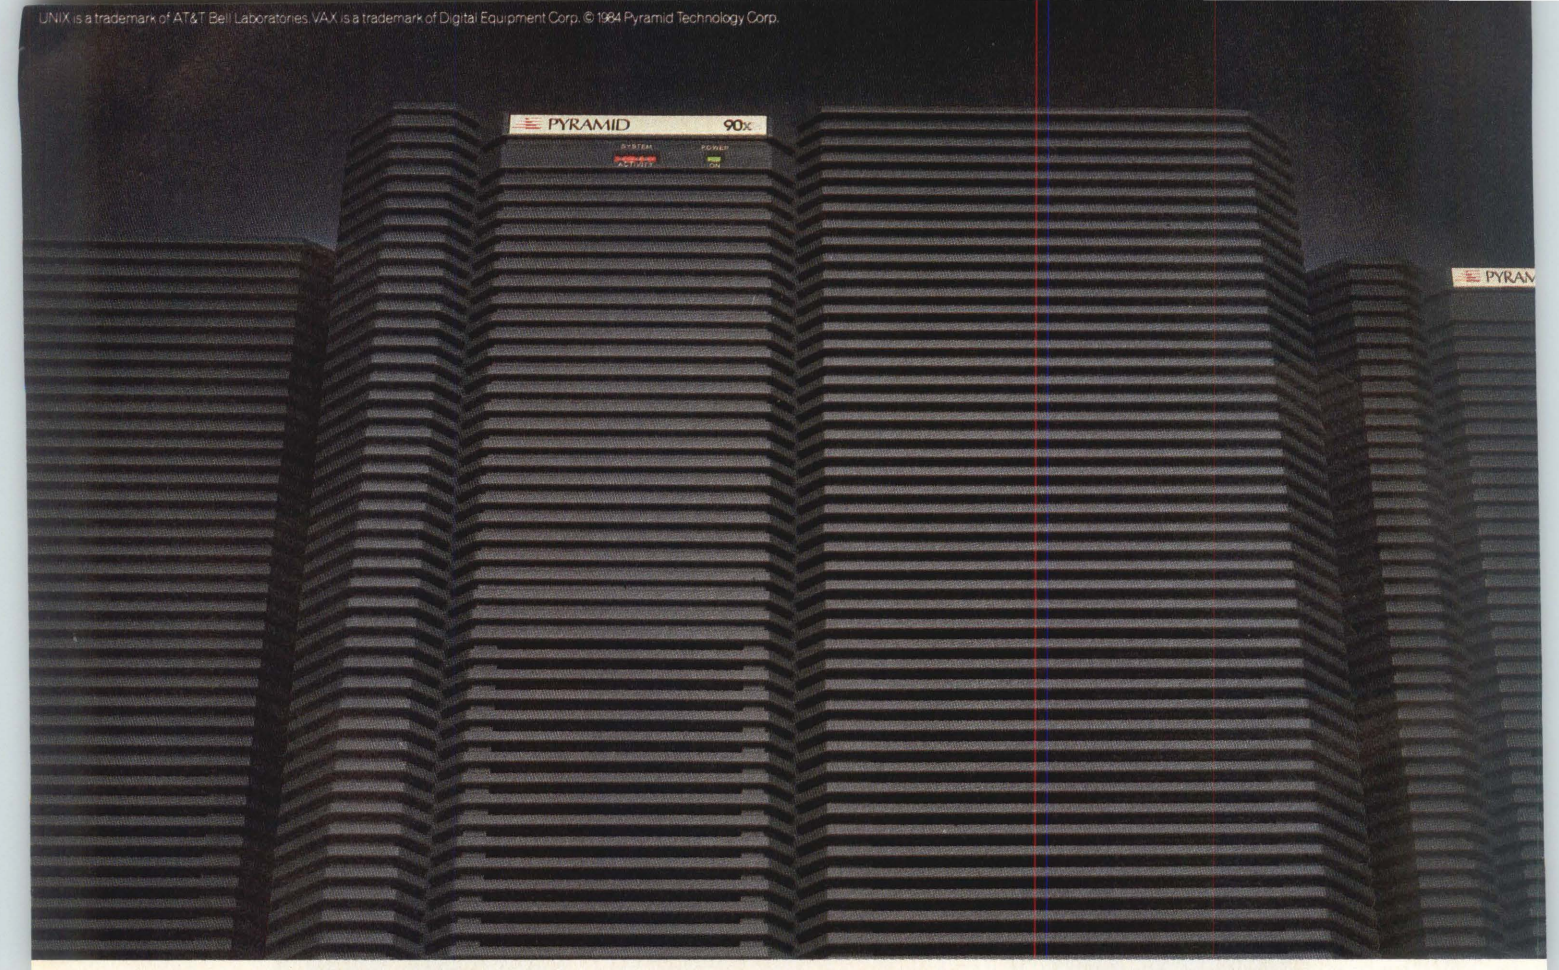
SYSTEM Z-from Z...ing

ZAX

Zax Corporation 2572 White Road, Irvine, California 92714
(714) 474-1170 • 800-421-0982 • TLX 183829

CIRCLE 67

UNIX is a trademark of AT&T Bell Laboratories. VAX is a trademark of Digital Equipment Corp. © 1984 Pyramid Technology Corp.



IF IT WEREN'T FOR UNIX WE NEVER COULD HAVE BUILT THE PYRAMIDS

Ordinary computers, yes. But not a Pyramid Technology 90x.

After all, here's a supermini not just capable of running UNIX,™ but born to run it. And run it up to four times faster than the most popular UNIX host. For a lot less money.

The secrets of this Pyramid are a thorough understanding of UNIX, a few fundamentals of RISC (Reduced Instruction Set Computer) theory, more registers than 30 VAXs, and a 32-bit proprietary architecture that outperforms a roomful of micros.

All combined to speed up UNIX just where it likes to slow down.

For example, gone are 85% of performance-robbing memory references. The endless parameter shuffling of yesterday's technology has been replaced with a hardware register window. Even context switching takes less than one percent of the CPU's time.

It's amazing what hardware architects can do, given the chance. It's almost as startling as what our software wizards did.

They crafted OSx, a dual port of Berkeley's 4.2 BSD and Bell's System V. Because you can switch environments at will, no UNIX port offers more capabilities. With absolutely no loss of compatibility.

Well, almost.

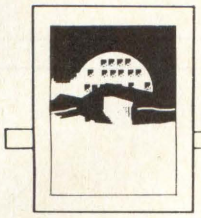
We do admit to one feature not compatible with other UNIX systems. Our single-source support.

One telephone number instantly connects you to both hardware and software experts. In-house pros, who spend their energy pointing you towards solutions. Not pointing fingers at each other.

So no matter how you see your requirements shaping up, contact Pyramid Technology, 1295 Charleston Road, Mountain View, California 94043. Or call (415) 965-7200.

Because when it comes to running UNIX, a Pyramid looks good from any angle.

PYRAMID TECHNOLOGY



JOINING TEXT AND GRAPHICS ENHANCES VIDEO PERFORMANCE

A dual-port RAM with a built-in shift register eliminates bottlenecks and speeds data transfers.

by David W. Gulley

Bit-mapped video graphics systems exemplify the need for higher density and higher performance semiconductor memories. Yet, all too often, these same memory devices are the bane of the system. The newest dynamic RAM devices, however, are allowing changes to the video graphics system organization. Thus, they are eliminating redundant support logic circuitry and providing a flexible system environment.

DRAMs, long associated with the frame buffer within the graphics section of a video system, provide the highest density and lowest cost storage for memory-intensive displays. High resolution graphics systems, such as those used in engineering workstations and computer aided design/computer aided manufacturing (CAD/CAM) terminals, require multiple memory planes to achieve the color capability necessary for a good user interface. In such a system, many parameters influence the available features while keeping the size and cost reasonable.

Often, the video display system designer is forced into "make-do" solutions when deciding on value-added features, especially where display memory is involved. Some features are common to many designs, and directly relate to the acceptance of a design in the market. Features considered high priority are the efficient integration of text and graphics, the time to redraw the screen image, the time to move objects onscreen, the amount of memory to map the display, and the support logic to use the memory effectively.

A typical video system contains separate text and graphics controllers (Fig 1). Thus, the system processor does not have to manipulate both the text display list and the graphics bit-mapped image. This system has evolved from the earliest text-only terminals, where there were no graphics requirements. In early systems, display memory consisted of perhaps 2 Kbytes for the display list RAM and 2 Kbytes for the character ROM. The need to place graphics images onscreen was first addressed using character graphics. By deepening the character ROM or adding a RAM to the character-generation circuit, user-defined characters could be produced.

To achieve more flexibility in image control, a bit-mapped memory is added into which the system can directly store images to be displayed. The mixed text

David W. Gulley is manager of MOS memory systems engineering at Texas Instruments, PO Box 1443, MS690, Houston, TX 77001. He holds a BS in electrical engineering from the Georgia Institute of Technology.

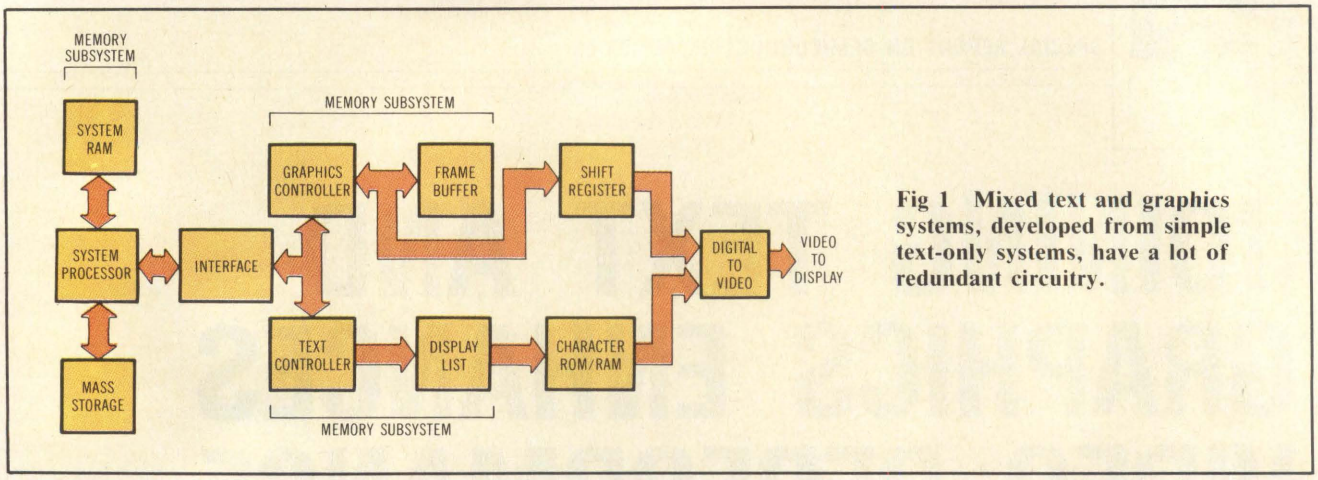


Fig 1 Mixed text and graphics systems, developed from simple text-only systems, have a lot of redundant circuitry.

and graphics solution is really a patch to add graphics capability to table-driven systems. However, future system design will treat text and graphics uniformly. New memory architectures are needed to make the transition to this type of system environment. The TMS4161 multiport video RAM is one device able to ensure this by providing a design path to the development of unified bit-mapped text and graphics systems (Fig 2).

Tracking the growth of video displays

Currently, uniformity is not in general use. Video display evolution has moved in another direction. As higher resolution and multiple gray-level or color planes were added, the screen refresh required higher data rates from the memory, giving less time to the system processor for data management in the frame buffer. As the resolution (pixels/in.^{2}) of the display increases by a factor of 2, the size of the display memory increases by 4, and the display interval for each pixel is reduced by a factor of 4.

The availability of dense, low cost DRAMs allowed expansion to higher resolutions (from a memory chip cost standpoint), but the DRAM architecture (1 bit wide) increased the data bus traffic needed to refresh the screen image. Graphics system controllers were added to the system to isolate the large bandwidth display bus from the system bus. If this isolation had not occurred, system processor throughput would have been seriously degraded. The data bus would be clogged with data passing from the frame buffer to the display.

The memory required for the frame buffer RAM is typically 10 (for black and white) to 40 (for 4 bits per pixel) times larger than the display list RAM in the mixed text system. Display data transferred to the screen loads the data bus so that there is considerably less time available to update the frame buffer memory than the display list RAM. Yet, since the nature of the data is single pixels, it requires more manipulating than display characters. More memory must then be accessed more often, and in less time. Hardware additions often implement many basic display functions, since there are not enough available memory accesses for software to optimally update the RAM.

The mixed video system consists of three memory subsystems, each containing a memory controller, memory logic, and glue logic. Glue logic also connects the controller to the system processor, and provides the required memory array drive. Each subsystem contains similar logic functions. Yet, the functions cannot be shared and still survive the data transfer bottleneck to the screen. Therefore, this is where DRAM features (actually, lack of features) have most influenced video system design. The many design approaches involving dedicated hardware control compensate for the limited accesses available to the memory. These approaches have partially relieved bus contention problems. But, the cost has

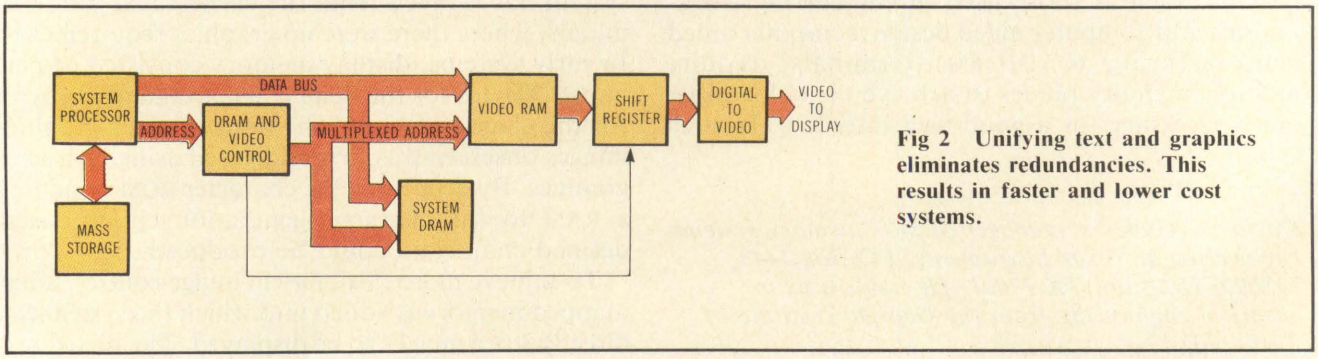


Fig 2 Unifying text and graphics eliminates redundancies. This results in faster and lower cost systems.

been a loss of system flexibility and compatibility for effective system upgrade. Dedicated controllers tend to lock the system into a set of fixed commands, character fonts, and data structures.

A high resolution (1024 x 1024 or 1280 x 1024) graphics display, as used in CAD systems, requires data from the refresh buffer at between 75 and 125 MHz from each plane, dependent upon the actual display device (monitor) specification. This is independent of the graphics controller's need to access the refresh buffer in order to update the image stored in memory. In the following analysis of system performance, a 1024 x 1024 noninterlaced display is used as a guide. Table 1 values describe the timings used in the analysis. Total frame time in Fig 3 consists of the active display interval, horizontal blanking interval (horizontal retrace), and the vertical blanking interval (vertical retrace).

TABLE 1
Display Parameters

Pixel clock frequency	88.00	MHz
Pixels per scan line	1380	MHz
Lines per frame	1063	MHz
Displayed pixels per scan line	1024	MHz
Displayed lines	1024	MHz
Horizontal blanking interval	4.05	μs
Vertical blanking interval	611.60	μs
Pixel time	11.36	ns

A typical video system design

The 88-MHz pixel data rate is in direct conflict with the need to update the memory quickly. The display refresh and the memory update must share the same data bus in the mixed text system. Updating the high resolution screen in a reasonable time frame requires some cycles to be available during the active display interval. A 1024 x 1024 display could be built using sixteen 64-Kbit DRAMs. But, even with the fastest parts, it is extremely difficult to get the video data rate required, and to be able to do useful screen image manipulations without reverting to a second (double) frame buffer.

The TMS4416 16-K x 4 RAM provides the large video bandwidth required in medium to high resolution video systems. Many systems that incorporate 16-K x 4 RAMs use the previous generation of 16-Kbit memories, and are using the x4 as a replacement for four 16-K parts. A wide-word architecture provides more data lines per depth of memory using standard DRAM access timing. Addressing four times as many bits per device simplifies the hardware needed to create the display frame buffer.

Wide-word devices used within the frame buffer provide the width needed to achieve the necessary bandwidth for display (Fig 4). This brute force design yields 64 data bits and requires a 64-bit shift register—all bits are loaded in parallel. The pixel clock is running at 88 MHz. S0 and S1 control the loading and shifting of the register. This approach contains the advantage of data access interleaving, first an interval for the processor access, and then an interval for the display access to the memory. It is more easily designed and manufactured than a similar approach using 16-K x 1 devices, and is much more reliable due to component and power reductions. The disadvantage is that there must be a way to buffer the data bus in order to convert from the 64-bit wide video section to the 16-bit wide system processor. In this design, a 64 to 16 multiplexer serves this function.

In the TMS4416 implementation of this circuit, there is one access available to the graphics controller for each display cycle. There are no highly critical access timings for the 16-K x 4, as the 64-bit shift register is loaded once each 727 ns (64 times 11.36 ns), and processor timing is assumed to be tightly coupled to the video shift rates. The storage cell refresh required by the DRAM is satisfied by reading across the memory chip rows for display accesses, and therefore does not require any additional logic or control. Even with all the data lines needed to connect the 64-bit shift register, this design runs at the top of its capability. If more flexible and higher performance systems are needed, the x4 RAM is not appropriate.

Many earlier high end video systems used the double buffer technique to avoid contention problems

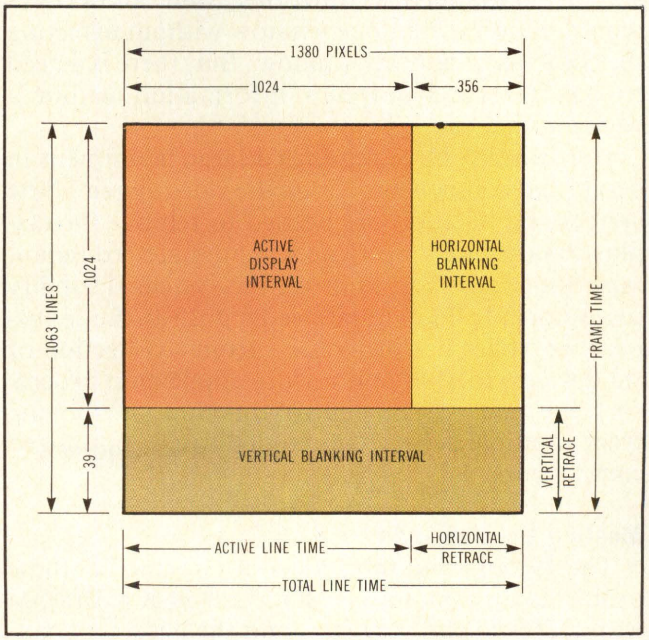


Fig 3 In a high resolution video display, the total frame time is the sum of active display time, and horizontal and vertical blanking intervals.

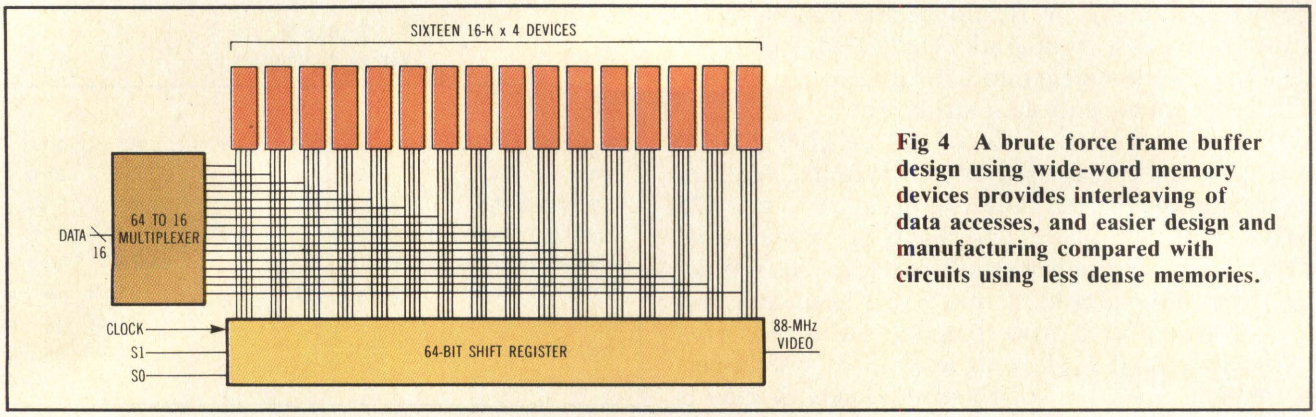


Fig 4 A brute force frame buffer design using wide-word memory devices provides interleaving of data accesses, and easier design and manufacturing compared with circuits using less dense memories.

between the graphics controller and the display refresh. In this scheme, two display frame buffers are used—one provides information for the display, the other is available to the graphics controller for updates. When the new drawing is complete, the system switches the function of the two buffers. Though this allows more interaction with the memory, it is at the expense of doubling the memory requirement. Also, when the buffers are switched, the graphics controller does not have a copy of the most recently available data image. In many systems, a form of DMA copies the data from one buffer to the other, effectively cutting the time available to the controller in half. Again, the system suffers from the lack of capability within DRAMs.

New systems are designed to be as functional to the end user as possible. The system must be flexible, tailored to individual needs, compatible with systems currently in use, and cost effective. Most new systems support multiple windows in order to display several simultaneous functions, and allow data manipulation within one window without affecting the contents of another window. But, there is a need to mix text and graphics information within a window.

Hardware control requires a large investment in design and components within the video system. New system architectures are needed to remove the display data-transfer bottleneck, eliminate redundant logic functions, and improve the system flexibility to conform to individual needs. Just as the industrial controller has progressed from a collection of SSI devices, to MSI, and now to single-chip processors, the video system control functions are moving from multiple subsystems to dedicated, optimized components.

Meeting the demands

The TMS4161 multiport video RAM remedies these problems by combining a standard 64-K x 1 DRAM with a 256-bit shift register, and the necessary controls to transfer data between the memory array and the shift register in a single package. By allowing simultaneous, asynchronous access to the two ports, the video RAM allows the system processor and the display refresh to work independently. Thus, the need for double buffering is removed, giving maximum time for the system processor to access the memory. The memory array access of the video RAM conforms to the signal and timing requirements of a standard DRAM. The onchip shift register supports high resolution data rates, and reduces video data shift logic and timing generation circuitry complexity. The shift register is configured as 4 linked 64-bit shift registers, able to provide shift lengths of 64, 128, 192 or 256 bits. These features help meet primary design criteria, and yield enhanced features for the video system.

The video RAM allows a more flexible approach to a video system design that eliminates mixed graphics approach patches. With the latter design, a bottleneck restricts data flow due to the single random access port on standard DRAM devices. Merging memory subsystems in a unified system substantially reduces design effort and cost. Redundant logic is eliminated by using the same functions for the video RAM as for system memory control. The logic required for the DRAM and video control section is currently implemented using several programmable logic arrays and MSI circuits (Fig 5). These could be placed in a gate array or other custom device.

The divider circuit is the only high speed device required, other than the external shift register, and provides the other logic with the appropriate timing signals. Not shown is the control to the external shift register, since it changes with implementation. A microprocessor or other controller can access the memory by issuing a MEMREQ/ with the appropriate read or write strobe. All other functions and timings are performed by the state sequencer.

A frame buffer using the video RAM could use a scan-line mapping architecture. This approach could also be used in the frame buffer of an existing design, although the full advantages of the dual-port would not be realized.

Scan-line mapping refers to positioning the memory devices to correspond to relative bit placements within a display scan line. Logic reduction in

Introducing Lear Siegler's 3278 Keyboard Compatibles

Easy to Look at. Easier to Use.

Now you can combine Lear Siegler's exclusive High Touch™ style with the convenience of true IBM 3278 keyboard compatibility.

The ADM 1178 video display terminals offer superior performance and ergonomic design. With standard protocol converters, you can interface with virtually any IBM mainframe to achieve substantial savings in hardware and operator training expense.

The ADM 1178 terminals can handle computer transmissions up to 19,200 baud without handshaking. They feature five video attributes (underlining, blink, blank, bold and reverse video), the IBM extended character set, four cursor modes (block or underline, blinking or steady), and 24 Program Function (PF) keys and two Program Access (PA) keys.

For operator convenience, the ADM 1178 terminals come with a full tilt and swivel monitor that stops positively in any position, an easy-to-read non-glare screen, and a detached, low-profile DIN standard keyboard.

The ADM 1178s are available with a standard 12″ green or amber screen and an RS-232C serial printer port. They can be easily modified for OEM applications and are available with such options as 14″ green or amber screen, answerback memory, current loop or RS-422 interface, and international character sets.

These Lear Siegler High Touch terminals are made in America-designed, engineered, manufactured and shipped from Anaheim, California. With this total, on-shore capability, and a complete worldwide network of sales and service centers, OEMs as well as end users can be assured of the best local support available in the industry.

Call our ADM 1178 product specialists today for complete information on products and protocol converters.



LEAR SIEGLER, INC.
DATA PRODUCTS DIVISION
901 E. Ball Road, Anaheim, CA 92805 (714) 774-1010

IBM is a registered trademark of International Business Machines Corp.

©1984 Lear Siegler, Inc.

CIRCLE 68

the frame buffer is evident, since there is only a 16-bit shift register, and no data bus buffer/separator requirement, as in the x4 example. In this particular example, each transfer from the memory array to the shift register moves a total of 4096 bits, which provide the data for four 1024-pixel scan lines. Onchip shift register data is loaded into the 16-bit shift register to accelerate the data to the required 88 MHz. The data in the memory's shift register is clocked at 5.5 MHz, well below the device's maximum clock frequency of 25 MHz. The timing for the video RAM is derived from the pixel clock to keep the system timings synchronous.

For this design, the row address strobe (RAS) cycle consists of 10 pixel clocks for the 114-ns precharge period, and 15 clocks for the 170-ns RAS low time, for a total period of 284 ns. All cycles (refresh, read, write, and transfer between arrays and shift register) use the same timings, with differences in the sequencing of the other control inputs to the video RAM (CAS/, W/, and TR/QE/). Each device holds every sixteenth pixel along the scan line of 1024 pixels. Scan-line data comes from 64 adjacent columns in each of the 16 devices. A 16-bit processor can directly access the memory array for image manipulation if it recognizes the appropriate addressing arrangement. Thus, the system processor can issue the address of the row and column for the desired pixel. The decoding of the active chip (when accessing via the DRAM port) may be done in hardware or as an internal operation of the processor.

The 256-bit register on the video RAM can be used by the video control logic to manipulate data as well as shift the data to the display. One way to employ this register is to clear (erase) the display quickly. The processor can write to the 256 locations corresponding to one row in the memory. This row can be transferred to the shift register. The shift register to memory transfer of the memory clears the remaining rows of the memory in 255 cycles. [Alternately, the serial input (SIN), could be grounded and SCLK clocked 256 times to load the shift register with all 0s.] Thus, the frame can be erased in a fraction of the vertical retrace interval of 612 μs, for improved performance in those applications requiring rapid screen clear.

Unlimited access

Since the video RAM shift register can be loaded from memory as little as once each four scan-line times for CRT refresh, the system processor has virtually unlimited access to the display memory. During a single 16.67-ms frame time, there would need to be 256 display access cycles (one of the video RAM's shift registers loads from memory for each four scan lines), and 1087 memory cell refresh cycles (a minimum of 256 refresh each 4 ms), which remove a small portion of the available time for updating

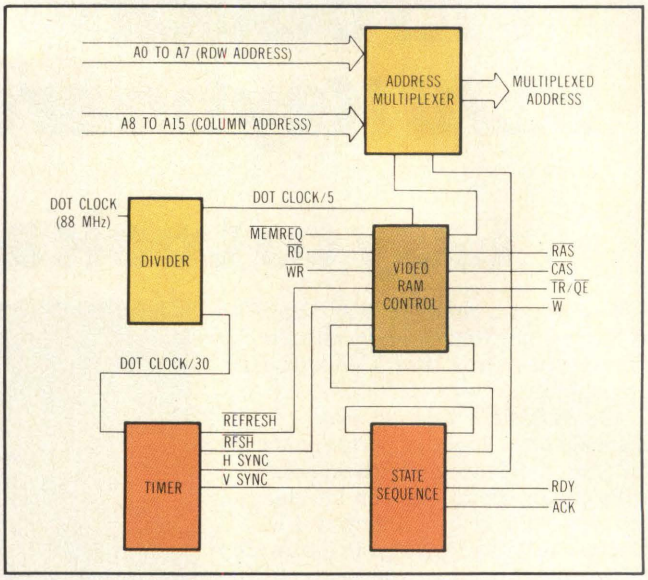


Fig 5 RAM and video control logic for a unified design can be implemented with programmable logic arrays and MSI circuits. A gate array or custom device could integrate the entire function in a high volume application.

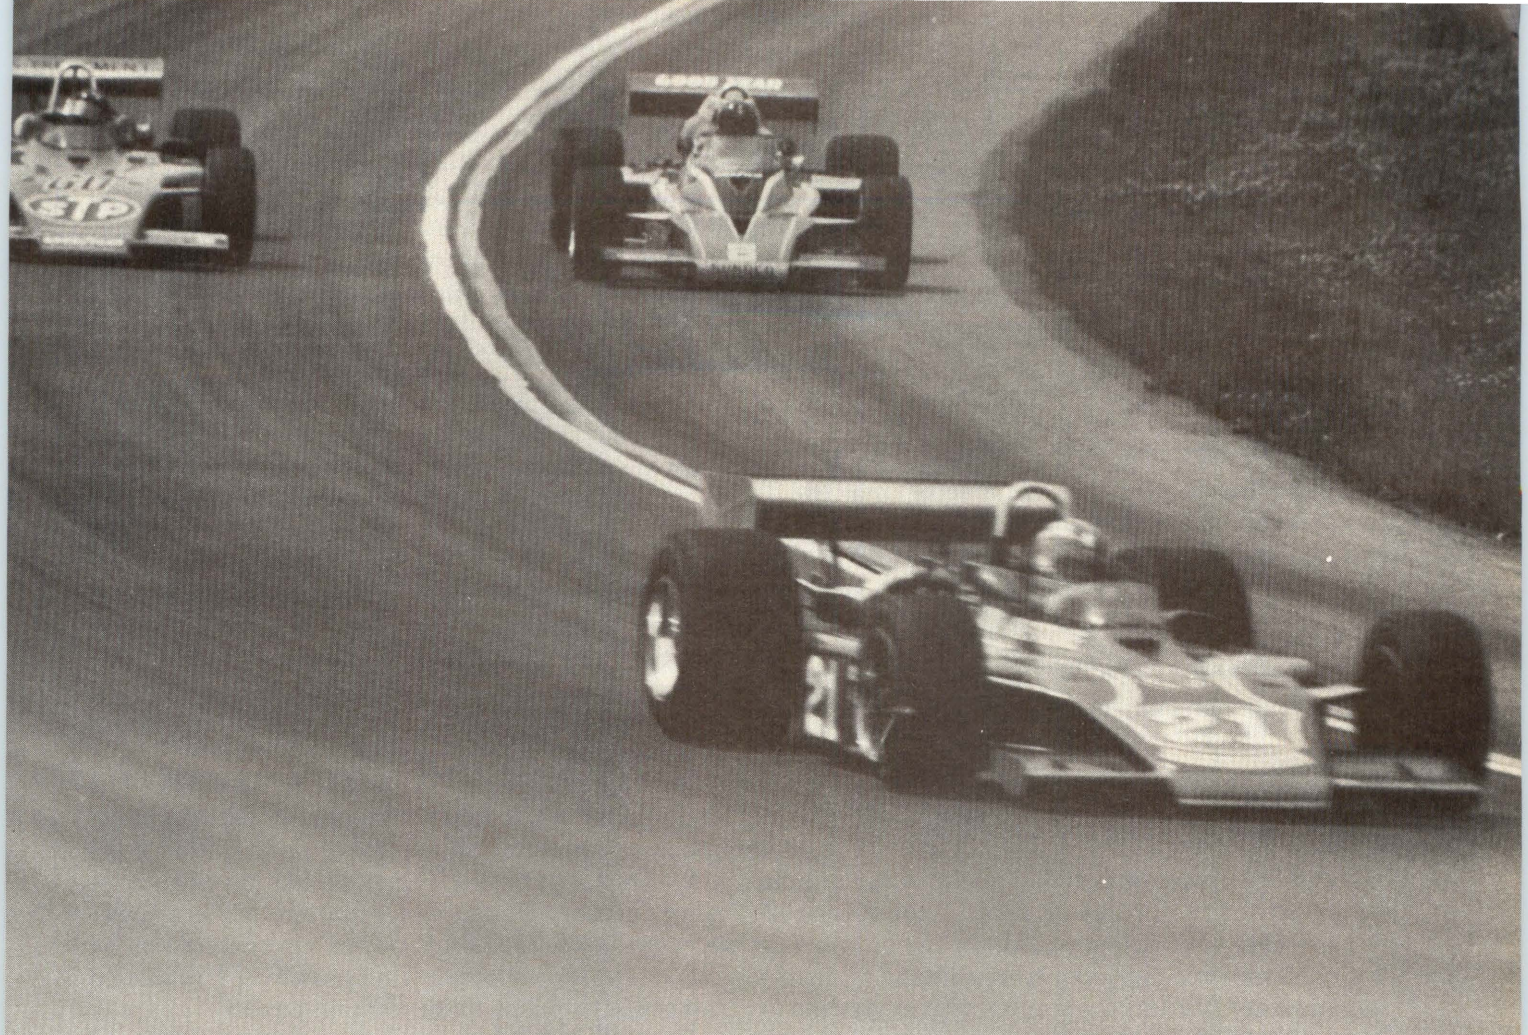
the screen. The time for this overhead can therefore be calculated as:

MC x (#DIS + #REF)
284 x (256 + 1087) = 381 μs

Where MC is memory cycle time,
#DIS is the number of display cycles, and
#REF is the number of refresh cycles.

So, in a single frame, all but 381 μs (about 2.3 percent) of the interval to be used by the system processor for display update are available. The remaining 97.7 percent of the time to scan a complete frame is available for access by the system processor. This allows memory accesses to follow logical, predictable patterns, and consistent timing sequences. These uniform cycles reduce the system hardware burden to fit memory update accesses into a narrow window or burst.

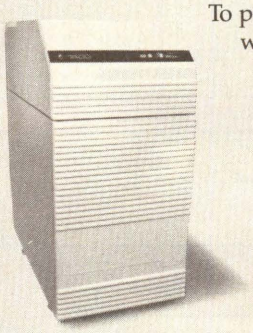
Modern screen imaging techniques indicate that hardware should not be used for read scrolling, to maintain maximum system flexibility. Designs usually call for moving data within defined regions of the frame buffer. However, for systems with hardware scrolling, the 256-bit register on the video RAM can be used by the controller to manipulate scan-line data in the displayed image. Data from one memory row can be transferred to the shift register, and then transferred back to another row (without shifting the data), which moves the pixel data from one displayed row to another. Several such transfers can be made, giving the effect of scrolling a full screen image vertically. This will scroll the entire width of the screen, so it may not be appropriate in a system with windows, where the scroll must be done in software.



Put your computer system back on the fast track

IBIS puts your computer system on the fast track. If you are building systems that require large-scale data storage *and* don't want to wait for data input/output delays, IBIS has the answer. Theirs is a new standard in disk drive technology. Especially if you are involved with image processing, super-mini systems, or other high performance computer requirements.

IBIS gives you 1.4 gigabytes of data storage on a single 14-inch Winchester disk drive. The data transfer rate is 12 megabytes per second with a track-to-track access time of 2.5 milliseconds and an average access time of 16 milliseconds. IBIS gives you the fastest drive in the industry with throughput that puts your system on the fast track.



There's more to an IBIS disk drive system than large-scale storage and a high-speed transfer rate. Our "designed-in" reliability assures your system's performance. IBIS utilizes the latest technology in thin film plated media and advanced air filtration. IBIS disk drives also have self-diagnostics, dynamic error detection, and modular construction that promotes ease of maintenance and service. In fact, with the high reliability of the IBIS drive, preventive maintenance consists of merely changing the pre-air filter once a year.

To put your computer system on the fast track with a reliable, high-performance disk drive, contact IBIS Systems, Inc., 5775 Lindero Canyon, Westlake Village, CA 91362. Telephone (818) 706-2505.

ibis™

CIRCLE 69

TABLE 2
Maximum Accesses to Arbitrarily Located Region

Region Size	Scan Line	Symmetric	Improvement
32 x 32	96	81	15 percent
18 x 18	54	25	54 percent
16 x 16	32	25	22 percent
10 x 10	20	9	55 percent
8 x 8	16	9	44 percent
4 x 4	8	4	100 percent

In scan-line architecture, a move of one row to an adjacent row within the video RAM results in moving the displayed line four scan lines vertically. To scroll an entire screen of lines would take:

(MTS + TC) x #ROW

(284 + 284) x 256 = 145 μs

Where MTS is the time for a memory to shift register transfer,

TC is the shift register to memory cycle time, and

#ROW is the number of rows to be moved.

This scroll operation could wrap the image around the screen, or the processor could update the display memory with a new portion of the image. The ability to move rapidly the screen image vertically may have application for some realtime systems or forms of animation, since it gives the system processor more time to update the displayed image.

The scan-line technique is preferred because it is simple and logical, and offers direct processor to memory mapping. Although scan-line mapping is generally best, other memory chip to pixel mapping schemes can be advantageous. In such a system, the drawing hardware may be able to update multiple pixels at each memory access. Unfortunately, it is exceptional for multiple pixels to occur in horizontal lines, such that writes could occur parallel to (along) the scan line. Data manipulations of the display involve the equally probable writing of multiple pixels vertically, diagonally, and horizontally to create an image. Most data manipulations involve pixels within an arbitrary region occupying multiple scan lines. Using the scan-line mapping technique, these arbitrary regions will most likely not align with the work boundaries accessed by the graphics controller, thus requiring multiple accesses.

Opting for the symmetrical architecture

One method to reduce the number of accesses necessary to transfer an arbitrary block, and to minimize the access of unnecessary pixels, is to use a symmetrical architecture for the frame buffer memory array. The symmetrical architecture uses one 4-bit shift register per plane rather than the 16-bit shift register of the scan-line approach. It cascades video RAMs by connecting the serial output of one device to the serial input of another to move 1024-column data bits to the 4-bit shift register. As in the scan-line method, an array-to-shift register transfer occurs once each four scan lines, but the data is now shifted out of the video RAMs at 22 MHz. The mapping to the screen shows that when the system processor operates on the frame buffer, it accesses a 4 x 4 block of pixels. The manipulation of an arbitrary memory image will generally require fewer accesses, since the number of pixels operated on by the system processor will be maximized (included unnecessary pixels will be minimized).

Table 2 compares the maximum number of accesses required to read or write variously sized, arbitrarily located regions using the implementations described. For the smaller regions, symmetrical mapping yields the greatest improvement in required accesses. The pixels of no interest occur at the boundary edge of the region. In addition, accesses internal to a large region do not contain any unnecessary pixels.

The symmetrical mapping architecture causes the scan-line data to correspond to the 256 columns within the same row of four memory devices. Each device corresponds to every fourth pixel in the scan

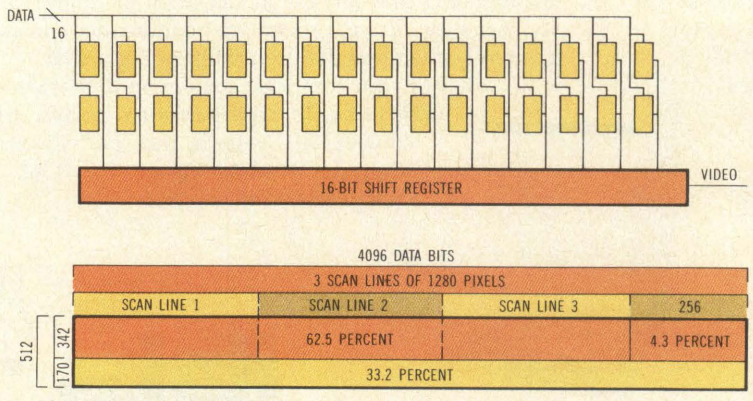


Fig 6 The video RAM organization offers the consolidation of memory, made possible by unified text and graphics design.

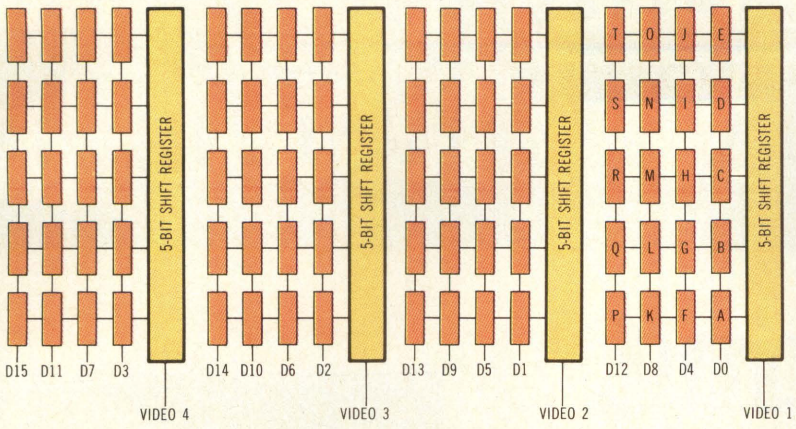


Fig 7 This frame buffer design provides four planes of display memory accessible to a 16-bit processor using 80 video RAMs.

line. A 16-bit processor can directly access the memory array for image manipulation, using the appropriate addressing arrangement. The system processor can directly issue the address of the row and column for the desired pixel. The decoding of the active chip may be done in hardware or as an internal operation of the processor.

To implement a 1280 x 1024 display (which is becoming somewhat standard), twenty 64-Kbit memory devices are required. There are apparent problems, however, with the use of a 16-bit processor with 20 memory devices. A bit-slice processor with 20 data bits could be used, but may not be practical for many systems. If a 16-bit processor is used, either the processor will access some or all of the memory as partial words (eg, 5 banks of 4 bits, or 1 bank of 16 bits and 1 bank of 4 bits), or extra memory is designated for use in video access. The use of partial words is possible. However, the added calculations to determine bit positions and increased number of accesses needed to update the display will cause some system performance degradation. This can be avoided by adding memory to fill out the data bus to a multiple of the processor width. This memory will not be wasted, since graphics systems typically require large regions of scratchpad memory to be used by the processors for placing text fonts, display lists, and for use in the calculations of drawing the displayed images.

Since more memory is required, the use of 32 video RAMs can simplify the task of matching the memory width to the processor width. If the memory is organized as shown in Fig 6, the transfer from array to shift register would place 4096 bits into the on-chip shift register. The data for 3 scan lines can be taken from these 4096 bits, leaving 256 unused bits. The display will use a total of 175,104 bytes (163,840 displayed and 11,264 left at the end of the rows) of the 262,144 bytes in the RAM. This noncontiguous memory amounts to about 4.3 percent of the total memory. The remaining 87,040 bytes consolidated within the second bank of video RAMs are available for use as system memory or scratchpad memory.

Lookup table eases calculations

To make the task of calculating the starting address of each scan line easier, a 1024-word table (2048 bytes), is set aside as a lookup table. Using a table to point to the start of the memory to be used for display allows rapid changes in portions of the screen image while not affecting other areas. When the same memory can be used for either display or system memory, the cost effectiveness and flexibility of the system is improved. The unified text and graphics design approach allows memory consolidation, especially in those systems where nonpower-of-two displays are used.

Fig 7 shows a possible implementation of a 1280 x 1024 frame buffer to provide four planes of display memory accessible to a 16-bit processor using 80 RAMs. The processor will access all four planes of data for each of four pixels, from what it considers as five banks of 16 memories. The data for display within each of the planes appears as four banks of five devices so that the array to shift register transfer will load four scan lines of data. The difference in this organization is the relative position of the four pixels accessed by the processor. The mapping separates the four pixels accessed by 1280 pixels into a vertical line. Depending on the address scrambling, the processor could map the memory sequentially in vertical rows rather than horizontal lines. The 5-bit shift registers allow the video dot rate to go up to 125 MHz before the data capacity of the RAMs is exceeded.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 710 *Average 711* *Low 712*

There Are Only Two Things New Tested Pairs

The Testing.

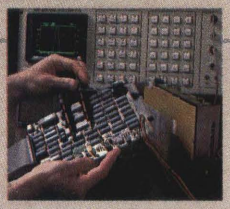




OEMs face major problems trying to decide which hard disk drives to build into their microcomputer systems. What with the wide range of drives, and so many hidden costs associated with evaluation, testing and integration, the process becomes a time-consuming and costly hassle.

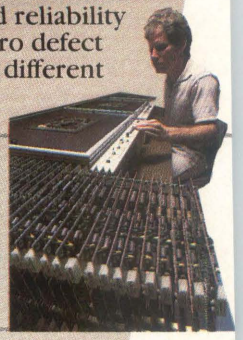
But there's a way out—from the company that's helped solve OEM disk controller integration problems more often and over a longer period of time than anybody in the business. That company is Xebec.

From standard-setting to vertical integration, from heavy investment in computer-aided design to the total commitment to the most automated manufacturing technologies, Xebec has demonstrated its leadership position in supplying microcomputer storage solutions.



Tested Pairs. The Guaranteed Match.

With its new tested pairs program, Xebec solves a major industry problem: post-delivery drive failure when interfaced to the controller. Having to do "after the fact" drive testing for many of our OEM customers, we have decided to offer a "before the fact" program. We'll guarantee quality and reliability by assuring a match between our zero defect controllers and a choice of drives in different capacities and form factors.



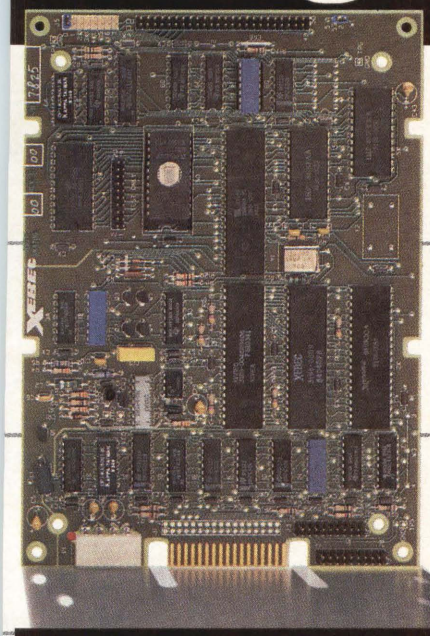
It's simple. Tell us which drives you're considering and we'll test them on Xebec-designed equipment against the most rigorous standards in the industry. Standards we at Xebec have set. Then we'll tell you which of our controllers is the best match for that drive—or we'll customize a controller for you—and accelerate your time to market.

The Edge In Controller Technology.

Xebec controllers are well known as the best in the industry. Their single-board designs incorporate MOS microprocessors and the latest standard cell and surface mount IC technology. Compatible with standard interfaces, our controllers have set the pace with sophisticated data separation, advanced error detection and correction, hard-fault isolation, a high-level

CORPORATE HEADQUARTERS
2055 Gateway Place, Suite 600
San Jose, California 95110 U.S.A.
(408) 287-2700

That Distinguish Xebec's Storage Solutions.



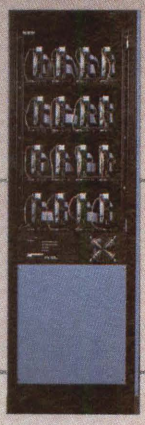
And The Pairing.

command set and other performance and reliability features.

Now these superior controllers can be paired with high-quality, Xebec-qualified drives. And you save time and money on evaluation engineering, incoming inspection, and service—while gaining a single-source, single purchase order solution to disk drive integration.

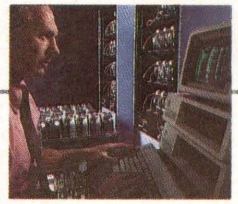
The Front Line Of Drive Testing.

We were recently asked to design and build a disk drive tester for one of the world's largest OEMs. We'll be using that tester to analyze disk drive performance for you. Full environmental testing, including shock and vibration. Full functional testing at elevated temperatures, with read/write tests at marginalized voltages. Careful calculation of hard and soft bit error rates. Complete checking of FCC and other agency emission standards. We'll make sure any drive meets its stated specifications and performance and quality claims.



How Xebec Does It. And Why.

Xebec can offer its tested pairs service to OEMs because of our commitment to quality. Zero defect quality symbolized by our "Xero D" signature and demonstrated by our superior computer-aided design, manufacturing and testing. Our talented people. Our experience. That's *how* we can provide high-quality products for prices no more than products that deliver a lot less.



Why offer our tested pairs? Why originate a new concept in quality? Enlightened self-interest. Because if we don't, someone else will. Frankly, we want to continue to enjoy the benefits of our leadership position as much as our customers do.

Call Xebec today. Pair up with one of our representatives to find out more about how Xebec tested pairs can benefit you.

Xero D.

XEBEC

The Zero Defect Company

SALES, INTERNATIONAL,
Belgium
32-02-762-9494
TWX: 65054 Xebec B
England
44-628-71138
TWX: 84443 Xebec G
Italy
39-6-344412
TWX: 622685

SALES, U.S.A.
Sunnyvale, CA (408) 733-4200
Irvine, CA (714) 851-1437
Atlanta, GA (404) 457-9872
Boston, MA (617) 740-1707
Dallas, TX (214) 361-0687
Baltimore, MD (301) 465-7771
Chicago, IL (312) 931-1420

U.S.A. AND CANADA DISTRIBUTORS
Kierulff
Hamilton-Avnet
Avnet Electronics
Hamilton Electro Sales

CIRCLE 70

READ. WRITE.

It is the highest performance, most reliable 5¼″ cartridge disk drive in the industry.

It features more resistance to shock and vibration than any other disk drive, fixed or removable.

Its cartridge is the least expensive among formatted 5-megabyte cartridges on the market today.

Its cartridge interchangeability from drive to drive is absolute. Its start/stop time is the fastest available of any high-performance disk drive.

It is the Beta 5 Cartridge Disk Drive from IOMEGA. And it is, in a word, superlative.

IOMEGA Corporation
1821 West 4000 South
Roy, Utah 84067
Or call (801) 776-7330



IOMEGA™

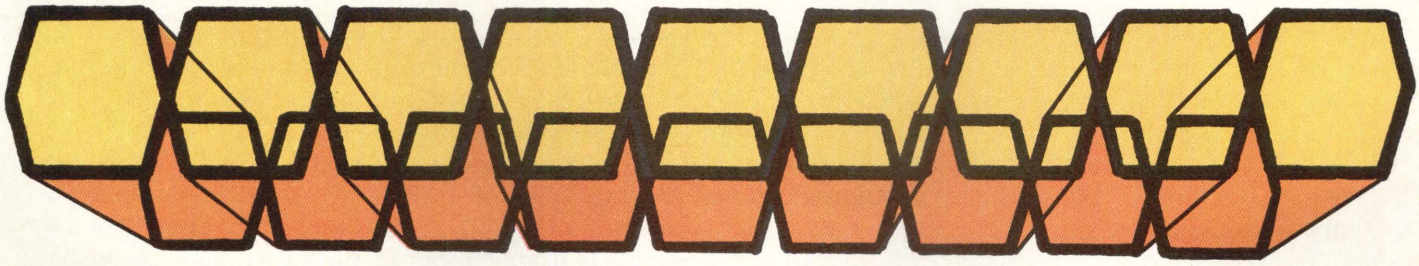
High Performance Cartridge Disk Drives

CIRCLE 71

REGIONAL AND DISTRICT SALES OFFICES: SOUTHEAST (305) 755-1060; SOUTHCENTRAL (214) 458-2534; WESTERN (714) 855-1211, (408) 263-4476; NORTHEAST (617) 933-2000; MIDWEST (414) 782-5229 EASTERN OFFICE (203) 359-9858.

CMOS 256-KBIT RAMs ARE FAST AND USE LESS POWER

By using a high performance CMOS dynamic RAM technology, relaxed timing margins combined with faster cycles can yield significant improvements in system speed.



by William H. Righter

Applications for dynamic RAMs continue to extend beyond mainframe computers, which have historically consumed the majority of such devices. Today, with the increased use of advanced microprocessor systems and increased sophistication of system software, there is a seemingly insatiable demand for memory. Thus, many of the new applications require memories with capabilities beyond those offered by standard NMOS DRAMs.

To provide these features, two new 256-Kbit DRAM product lines, part of Intel's family of CMOS DRAMs, make use of a high performance technology called CHMOS-D III. This technology combines the low power of CMOS with the performance of the HMOS III technology.

Since the new devices consume less power than their predecessors, they provide longer battery life for applications such as portable instruments and computers. The low power consumption of CHMOS technology makes it practical to add static column circuitry. When combined with a self-timed write function, this circuitry yields shorter device cycle times and wide timing windows within the cycle. This ability results in fast, usable speeds at the system level. When combined with Ripplemode, a fast access mode, static column decoding provides rapid and random access to all data bits in a device row. This large data bandwidth is useful in high performance graphics applications or array and signal processing systems.

The high density, 256-Kbit DRAMs also yield high board densities compared with other DRAMs or static RAMs. The 51C256 family, for example, is organized as 256-K x 1 (Fig 1) and is a density upgrade for the 51C64 CHMOS-D III DRAM family. It is packaged in a standard 16-pin high reliability plastic

William H. Righter is a senior technical marketing engineer at Intel Corp's Memory Products Division, 2111 NE 25th St, Hillsboro, OR 97123. Mr Righter is responsible for applications of CMOS dynamic RAMs. He has studied computer engineering at the University of Portland.

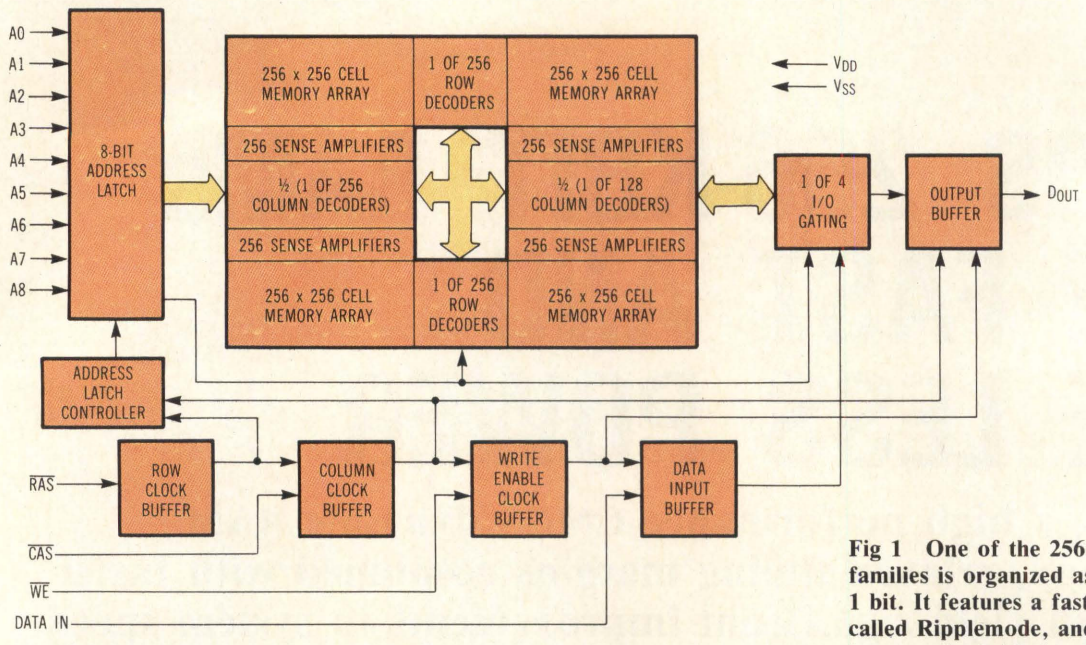


Fig 1 One of the 256-Kbit DRAM families is organized as 256-K x 1 bit. It features a fast access mode called Ripplemode, and is packaged in a standard 16-pin plastic DIP.

DIP and features the high data bandwidth operation of Ripplemode. The second new DRAM family, the 51C259, is organized as 64-K x 4 and is packaged in an industry standard 18-pin high reliability plastic DIP. This family features static column mode operation.

Attaining relaxed timing margins

The NMOS DRAMs present the problem of tight and restrictive timing relationships between various edges of input signals. These timing relations are so constrained that it is nearly impossible to design a high performance system that operates at minimum device cycle times. Moreover, if system skews just consume the timing windows on a 260-ns NMOS DRAM, they could literally overrun and cause a 65-ns CHMOS RAM (whose timing windows were scaled proportionately) to be useless. Relaxed timing margins, not just faster cycles, are needed for usable speed at the system level. Static column decoding and a self-timed write operation provide these relaxed timing margins.

The internal column address buffer in the 51C256 DRAM family is a flow-through latch. During the portion of the read cycle when the column address strobe (CAS) is high, addresses flow through the buffer. This means that the data access begins at the instant the column addresses are valid, instead of from the (later arriving) leading edge of CAS. This new access specification is called t_{CAA}. The design eliminates the CAS clock from the critical access path of the system. When CAS does become active, it latches the column addresses and turns on the output buffer. The system timing skews become more tolerable when tight address setup times to CAS are removed, and the need for tight control over the leading edge of CAS or t_{RCD}—the row address strobe (RAS) to CAS delay—is eliminated. The net result is usable speed in the system.

A memory cycle on the 64-K x 4-bit 51C259 device is somewhat different. Two major differences are found in the function of CAS and the inclusion of an output enable (OE) pin. On these devices, toggling CAS is optional because the internal column address buffer is permanently wired in flow-through mode (for read cycles). For this reason, the 51C259 can operate with CAS at V_{IL} indefinitely. Operation with CAS grounded is also possible. In this mode, once RAS has latched row addresses, the data access begins as soon as column addresses become valid. The access specification is again called t_{CAA}, and is measured from valid addresses to data output. Because data I/O occurs on the same pins, a fast response OE is provided so that bus contention on the data I/O lines can be avoided.

An essential design goal is to have the system run as closely as possible to minimum device cycle times. One of the limiting factors during write cycles involves controlling the width and position of the write enable (WE) pulse. After accounting for worst-case system skews, it may not be possible to generate a minimum write pulse width (t_{WP}). Yet, at the other extreme, if t_{WP} is too wide, the positioning of the trailing edge of WE will force the cycle

WE SOLVE PROBLEMS!



In all sizes, shapes, colors and configurations.

You need high performance, innovative design data displays. Contact Audiotronics. We have what you need. In 3″, 5″, 7″, 9″, 12″, 14″, 15″, 23″ and 5″ × 9″, monochrome and color, integrated (neck-mounted), chassis, kit or cabinet. And if we don't have what you need, we'll design it for you.

Using our basic displays, our engineers become your engineers. They custom design a display for your specific application, meeting your particular system design requirements.

Just give us your specifications and we'll solve all your display problems. And we'll deliver on time, when you want them.

Audiotronics has been solving problems for almost 30 years, designing thousands of custom data displays for important customers like you, both large and small. Call us today. Turn our engineers loose with your display system problems.

AUDIOTRONICS

North Hollywood
California 91605
(818) 781-6700

CIRCLE 72

time to be lengthened by the WE precharge or some other specification.

Particularly during the rapid cycle of other static column DRAMs, the small window provided between the minimum t_{WP} and the maximum (without pushing out the cycle) allowable system t_{WP} is absorbed in system skews. This forces a longer cycle time just to get a write cycle to work at the system level. To alleviate this difficulty, the new CHMOS 256-K DRAMs feature a self-timed write pulse that eliminates the need to tightly control the write pulse width and its trailing edge. This removes many restrictive WE timing requirements from the critical path. The benefit of the self-timed write is the same as removing CAS from the critical timing list. As a result, fast and usable speeds at the system level are more easily attained.

Because of the common I/O configuration of the 64-K x 4 51C259, the operation of a write cycle differs slightly from usual DRAM write cycles. For example, the concept of an early write does not apply to the device if it is operated with CAS at ground (V_{IL}). Early write cycle timing is referenced to the leading edge of CAS. Without a CAS transition, there is no specified early write in the usual sense. The purpose of early write cycles is to maintain the data output in a high impedance state so that data can be driven into the device I/O lines without causing bus contention. However, the same result is attained by maintaining the OE line high.

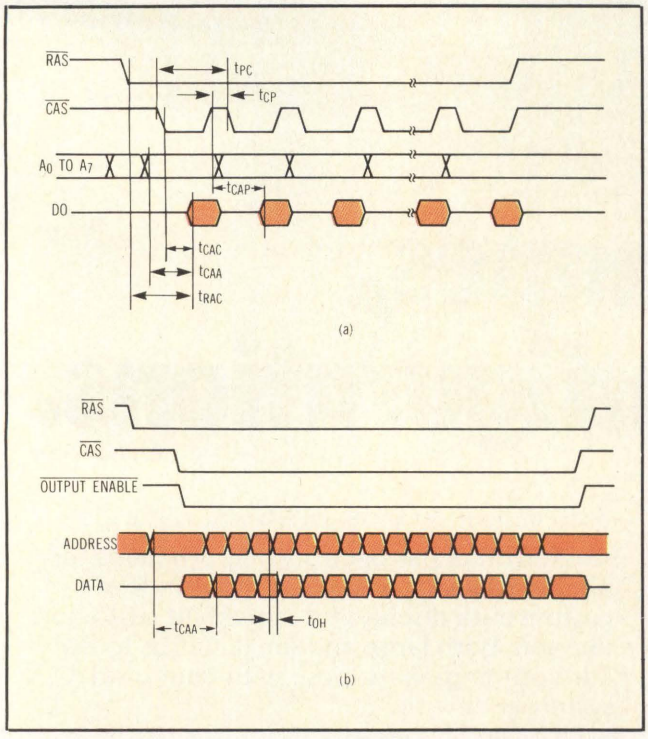


Fig 2 Addresses flow through the DRAM while column address strobe (CAS) is high, so data access actually begins from valid column addresses (t_{CAA}), not from CAS (a). Eliminating the CAS clock from the critical path of the DRAM makes short cycle times easier to use. In the static column mode (b), one clock period is eliminated, providing simpler system design. In this mode, the DRAM acts like a fast static device.

Late write cycles and read-modify-write (R-M-W) cycles are similar in that both are generated by bringing WE low sometime after CAS. However, for a R-M-W cycle, all of the minimum delays to WE must be met (t_{RWD}, t_{CWD}, and t_{AWD}). By satisfying these minimums, D_{OUT} has time to become valid. On the 51C259, the four I/O lines must be controlled with OE to avoid bus contention between the read data and the data to be written during the cycle.

Ripplemode cycles

The 51C256 DRAMs include Ripplemode, a high speed operating mode that permits all 512 data bits within a selected row to be randomly accessed at a high data rate. Ripplemode is compatible with page mode in other DRAMs, but has faster access and cycle times. The basic timing is shown in Fig 2 (a). Because addresses flow through the DRAM while CAS is high, the data access actually begins from valid column addresses (t_{CAA}), rather than from CAS. This technique eliminates the CAS clock from the critical path of the DRAM system, making the rapid cycle times easier to use.

Another access specification [Fig 2 (b)] is called t_{CAP}. This is the access time as measured from the column precharge, or the trailing edge of the previous cycle CAS, to the time of valid data in the current cycle. This is a necessary specification because, although the column address buffer has a flow-through latch, the latch does not open for flow-through operation until CAS returns high. Since addresses are latched by CAS, the addresses can usually be changed well before the end of the cycle, thus pipelining addresses for the next cycle. The limiting factor in this technique is the access time from the trailing edge of CAS, which is the earliest that the addresses can flow through the latch and begin to access data. To guarantee valid data, both t_{CAA} and t_{CAP} must be met.

Ripplemode cycles are entered with a normal RAS/CAS sequence. After this entry cycle, in which row addresses are latched, RAS is maintained low throughout the time period required for the fast column accesses, up to the maximum RAS low time for Ripplemode (t_{RPM}) of 75 μs. CAS is then continuously toggled at the required data rate. Minimum access time throughout this period is usually t_{CAP}. In this mode, the DRAM functions very much like a 512 x 1-bit high speed latched (or synchronous) static RAM.

Static column mode operation is the high data bandwidth mode of the 51C259 family. The basic timing is shown in Fig 2(b). The static column mode permits all the data that is within a device row

(256 addresses) to be randomly accessed 4 bits at a time. During static column mode operation, read, write, and R-M-W operations can be performed. The cycle is entered by activating RAS, which latches a row address. Then, while maintaining both RAS and CAS low, column addresses are rapidly cycled, selecting the 4 data bits that are directed to the I/O lines. This technique eliminates one clock (CAS) from the fast access mode, which provides a simpler system design. The 51C259 in static column mode behaves like a 256 x 4-bit fast SRAM in that each time the addresses are changed, 4 new data bits appear on the I/O lines within the specified address access time (t_{CAA}).

With the new t_{CAA} access specification and static column mode, it is sometimes unclear whether system access is limited by t_{RAC}, t_{CAA}, t_{CAC}, or t_{OE}. Access time formulas can be produced from data sheet specifications, but the easiest way is to draw a simple worst-case timing diagram that shows the relationship between RAS, addresses, and CAS as they occur in the system being analyzed. Measure the access time from each of the three edges based on the data sheet access specifications. Draw all three data output waveforms. The output data on the timing diagram farthest from RAS is the true system access time. Remember that although the data output becomes active just before the t_{CAC} time period expires (t_{LZ}, or t_{RLZ} in grounded CAS systems, to be exact), the output data is not valid until t_{RAC}, t_{CAA}, and t_{CAC} time periods have elapsed.

Graphics design

Static column architecture in the CHMOS DRAMs provides very high sustained data rates. Coupled with a symmetrical cell addressing technique, static column DRAMs provide a solution to requirements for high performance bit-mapped graphics memories.

There are two key requirements for memories used in high performance graphics systems. The first is to minimally meet or, preferably, exceed the high serial bandwidth required to refresh the display. The other key requirement is to maximize the rate at which the display can be updated, because this is the most visible sign of performance to the end user. CHMOS DRAMs can meet these requirements.

Required Memory Bandwidth as Inverse of Pixel Time

Display Size	Pixel Time	Bandwidth Required
512 x 384	66.5 ns	15 MHz
640 x 512	38.3 ns	26.1 MHz
1024 x 768	14.4 ns	69.4 MHz
1024 x 1024	10.0 ns	100 MHz
1280 x 1024	8.0 ns	125 MHz

Graphics requires a high speed serial stream of data.

A 1280-pixel x 1025-line graphics display provides an example.

A pixel is the smallest displayable screen image, and the pixel time is the amount of time allocated to display a given pixel. The required instantaneous memory bandwidth is simply the inverse of the pixel time. The Table shows the required memory bandwidths for various screen sizes. These bandwidths are usually achieved by parallel loading memory data into shift registers and shifting the bits out with the pixel clock. Fig 3 shows one way to construct one bit plane of a 1280-pixel x 1024-line display frame using 64-K x 4-bit CHMOS DRAMs. It takes five devices to provide enough bits for this size display. By arranging the memory devices in parallel, there are 20 display pixels available per memory cycle.

The display bandwidth is met by using static column mode cycles to burst-fill the video output registers (VORs). The VOR is constructed with 4 x 4 register files (for example, 74S670 devices), and is arranged as 5 devices wide (20 bits) x 2 devices deep (8 register layers). While the register files are being emptied, the memory is available for updating by the pixel engine. The static column mode ensures that the bit-mapped memory can exceed the display bandwidth requirements, providing enough time to steal cycles for random update in between burst fills of the output register. Cycle stealing is the technique of interweaving random update cycles with the sequential read cycles required for the display. This method requires a higher bandwidth memory, but provides high performance by allowing updates to occur during the display time without flicker.

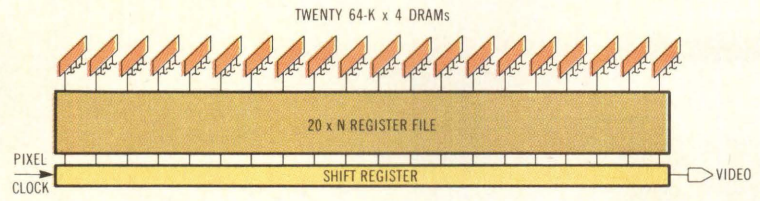


Fig 3 In a high performance bit-mapped graphics memory, it takes twenty 64-K x 4-bit DRAMs to construct four bit planes of a 1280-pixel x 1024-line display frame. Arranging the memory devices in parallel, 20 display pixels are available per memory cycle.

In a high performance frame buffer design, the ultimate limiting factor on the update rate is the memory cycle time, which for most NMOS DRAMs is around 260 ns. Ideally, one would like to take advantage of the 55- to 65-ns cycle time of static column mode to optimize the use of time available to update. A simple address mapping technique, called symmetrical cell addressing, allows the high speeds of static column mode to be used for drawing graphics objects.

The static column architecture provides high sustained data rates.

In a conventional linear-addressed frame buffer, sequential bits in a memory row are mapped to successive pixels on a display line. For a 1280 x 1024 display, this conventional method of addressing maps the bits from the first internal row of the memory devices into the first four display lines. With this organization, static column mode is not very useful for updates except for events occurring within the same line, such as drawing horizontal lines or clearing the screen.

If, instead, the device row addresses are mapped into a symmetrical cell on the display screen, then all random updates that occur within a given cell can occur at static column speeds, maximizing use of the time available for cycle stealing. Only when a cell boundary is crossed will a normal RAS/CAS cycle need to be executed. Thus, cycle stealing with static column decoding, coupled with symmetrical cell addressing, can dramatically improve update bandwidth.

In a 1280-pixel x 1024-line frame buffer, using static column mode CHMOS DRAMs with symmetrical cell addressing and a pixel clock of 8 ns (Fig 4), memory device cycle times are

Normal RAS/CAS cycle time:
32 pixel clocks = 256 ns
Static column cycle time:
9 pixel clocks = 72 ns

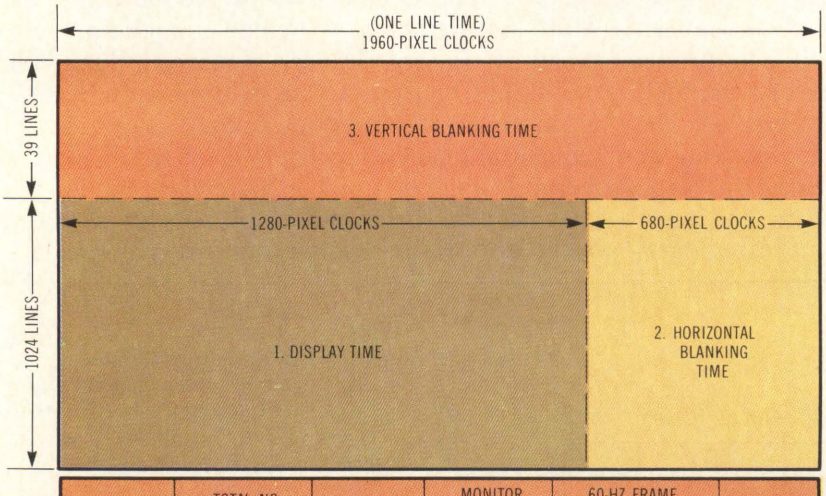
A four bit-plane 1280 by 1024 frame buffer requires twenty 256-K memory devices. This provides 5120 bits (20 parallel data bits times, 256 locations per device row in the 64-K x 4 DRAM) to map into the symmetrical cells. The cells are then 80 bits wide by 64 lines deep, reflecting the 5:4 aspect ratio of this display. The video output register (VOR) structure is 20 bits wide and eight registers deep to allow cycle stealing. To determine time available for cycle stealing:

Empty VOR	160 bits x 8 ns	=	1280 ns
Fill VOR	512 ns + 6 x 72 ns	=	− 944 ns
Time available to cycle steal		=	36 ns

To determine the maximum screen update rate, the average time to perform updates (write cycles) must be known. Assume that a randomly drawn vector (using Bresenham's algorithm) will have, on an average, 32 pixels updated within an 80- by 64-bit cell before crossing a cell boundary. The actual calculation accounts for one normal memory cycle time to enter static column mode (256 ns) plus 31 cycles at static column speed. The average static column cycle time is then:

(256 ns + 31 x 72 ns) /32 = 77.75 ns

With 336 ns available to cycle steal, an average of 4.3 updates can occur each 1280-ns interval, of which there are 8 per line. Multiplying by the number of



DISPLAY SIZE	TOTAL NO. OF PIXELS	PIXEL CLOCK	MONITOR		60-HZ FRAME		PIXEL TIME
			VB	HB	DISPLAY	BLANK	
1280 x 1024	1310-K	125 MHz	611.52 μs	5.44 μs	10.48 ms	6.1866 ms	8 ns

Fig 4 Memory device cycle times can be calculated for the 1280-pixel x 1024-line display frame. The normal RAS/CAS cycle time is 256 ns; static-column cycle time is 72 ns.

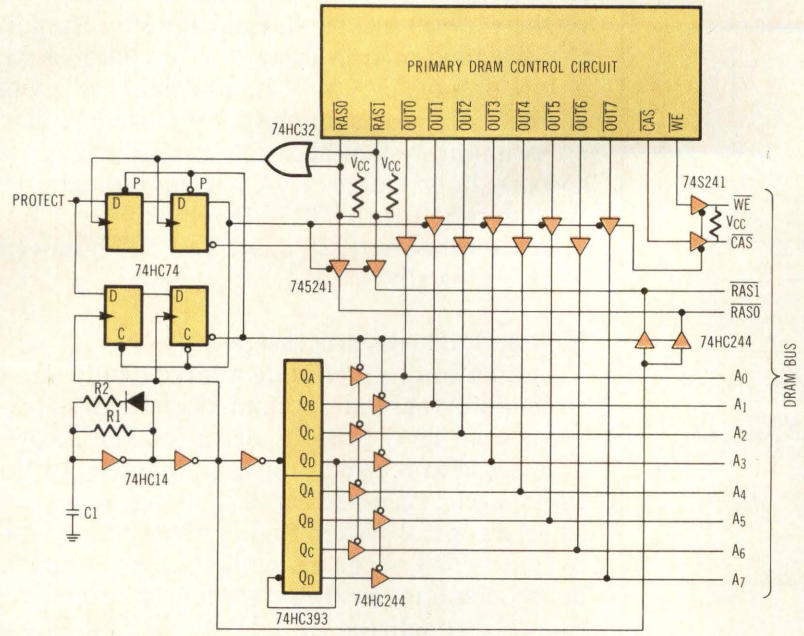


Fig 5 To illustrate the suitability of CHMOS DRAM for battery-backup memory systems, this circuit shows a standard VLSI design providing the main DRAM control and refresh during normal operation. Discrete CMOS logic is used to provide refresh during battery-backup operation.

display lines (1024) yields 35,225 possible updates during the display time. Adding in the number of updates during the blanking period gives the total number of possible updates each frame. The update bandwidth is N:

N = 60 x (35,225 + 6.04 ms/72 ns)
N = 7.14 million updates/s

The best possible performance for a standard or dual-port NMOS DRAM is about three million updates per second.

It is evident from these simple calculations that 64-K x 4 CHMOS DRAMs provide a very high performance graphics memory. They are a bandwidth match to the highest performance TTL or bit-slice pixel engines. In the near future, the trend to integrate more functions onto a single silicon chip will yield VLSI graphics controllers that can take advantage of this performance.

Portable systems

Another benefit provided by CHMOS DRAMs is low power consumption, which is particularly beneficial in battery-backup systems or in portable systems. In these applications, "L" versions of the CHMOS DRAMs (eg, the 51C256L or the 51C259L) can significantly extend the battery life of the system.

There are three major design rules for constructing low power systems. First, drive RAS to CMOS high levels (V_{DD} = 0.5V). To reap the lowest power benefits in other CMOS technologies, CMOS levels (not TTL levels) must be used to drive all inputs. On the CHMOS DRAM, driving RAS to CMOS levels (with CAS at V_{IH}) is the only requirement. The second rule is to take advantage of the extended refresh period of the "L" version to reduce the amount of time that refresh current is drawn during standby operation. The extended refresh cycles also lower the power consumed by the DRAM controller, especially if it is also CMOS. The third rule is to minimize the period that RAS remains low during any DRAM cycle. This minimizes the time that current higher than the standby current is drawn and reduces the amount of power consumed over any given period of time.

The circuit in Fig 5 is one example of how these three rules can be combined in a system designed for battery-backup operation. A standard VLSI (or TTL, or gate array, or programmable array logic) design provides the primary DRAM control and refresh during normal operation. When the primary system power is removed, the DRAM controller is powered down and isolated from the address bus. With the battery backup power provided, DRAM refresh is maintained by the low power discrete CMOS circuitry. This method is an effective alternative when the normal DRAM controller consumes a significant amount of power. For proper operation of this circuit, the signal, which typically originates at the power supply, must lead the primary system power going down (out of spec) by 100 μs, and must lag the system power supply on power-up by 100 μs after it is within the specified limits.

The CMOS refresh control circuitry is divided into three sections: the bus isolation circuitry, the sequencer that switches the two DRAM controllers in and out of the circuit, and the extended refresh circuitry. Bus isolation circuitry removes the CMOS refresh counters from the bus when primary power is applied, and isolates the primary DRAM controller

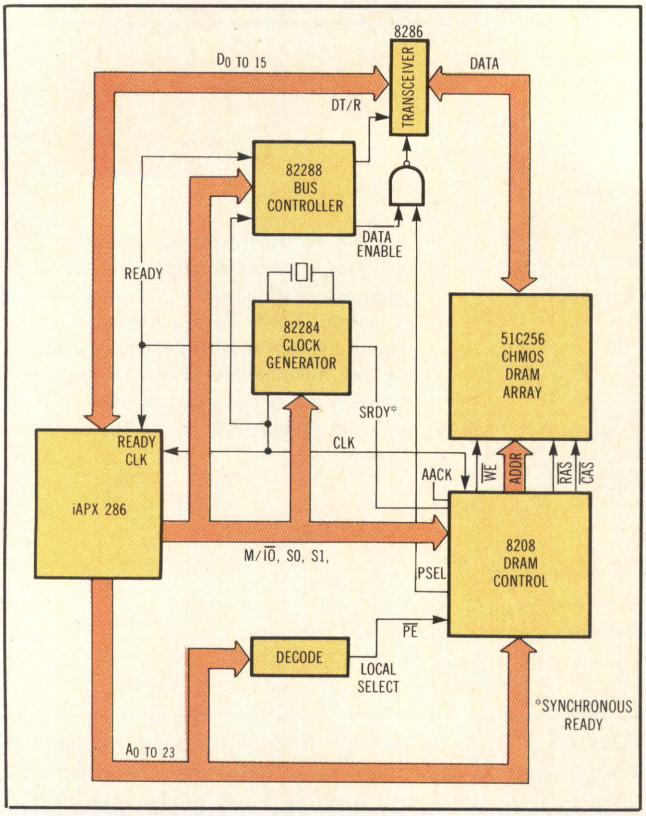


Fig 6 Fast cycle times and relaxed timing margins within the cycles are useful in high performance and high density memory systems.

from the address bus during battery-operated standby mode. The isolation circuitry for the primary DRAM controller consists of a 74S241 high speed Schottky three-state buffer, which is used for two reasons. First, it minimizes the timing skews in the address and control lines, allowing full-speed operation of the primary DRAM controller. Second, the high speed Schottky buffers, which lose power with the rest of the system, have open circuit outputs. Some other logic families have outputs that would drag the bus to either ground or V_{CC}.

The function of the sequencer is to switch control of the DRAM address bus between the CMOS refresh controller and the primary controller in an orderly fashion. It consists of four 74HC74 flipflops. Half of the sequencer is synchronized with the primary controller so that it is never switched on to or off of the DRAM bus in mid-cycle. This prevents violation of the minimum RAS low-time specification of the DRAM, which could imperil data integrity. For the same reason, the discrete refresh circuitry must not be switched on or off the DRAM bus during a refresh cycle, so the other half of the sequencer is synchronously switched with the refresh clock.

The discrete refresh circuitry consists of an RC oscillator and a binary refresh address counter. The counter is incremented on the rising edge of RAS. The duty cycle of the oscillator is designed to provide a minimal RAS low time, thereby minimizing the DRAM current drain. The frequency of the oscillator is set to provide the extended refresh interval of the CHMOS DRAMs during standby. Because of the extended refresh cycle, the oscillator is run at a low frequency, thus minimizing the dynamic current component of the total current drain.

Assuming system backup power is provided by a 450 mA-hour battery (AA NiCads), the minimum backup time for this 0.5-Mbyte memory would be about 150 hours (6.25 days) and 500 hours (three weeks) typically.

High performance processors

In addition to providing a large bandwidth data stream, internal static column decoders also provide the system level benefits of fast single cycles and accesses with relaxed timing parameters within the DRAM cycle. As noted earlier, it is generally not possible to control system timing skews tightly enough over a broad range of temperature, propagation delays, loading, etc, to achieve the minimum device timing through the critical access path. Meeting t_{RCD} (the RAS to CAS delay time) and the setup and hold times of addresses to CAS (which include skews in the CAS clock and addresses) can prove troublesome. These system skews are directly in the access path, and so contribute to a lengthening of the system access time. CHMOS DRAMs with static column decoding eliminate these tight timing difficulties by removing CAS and associated timings from the critical path. This provides much wider timing windows to absorb skews, so that high performance system design is easier to achieve.

The speedup in cycle times and the relaxed timing margins within the cycle are very useful in designing high performance and high density memory systems. Fig 6 depicts an example of such a system. In it, the 256-K x 1 DRAMs make up a high speed local memory for a 6-MHz iAPX 286 microprocessor system. The memory is accessed under control of an 8202 DRAM controller, using a synchronous status interface to the iAPX 286. The access time of the DRAMs is determined by the timing edge of CAS provided by the 8208 and hence is limited by t_{CAC}. In addition to the fast t_{CAC} requirement of this system, the overall access path requires a DRAM with a t_{RAS} of 150 ns. CHMOS DRAMs can meet these requirements and provide no wait states at 6 MHz. The 256-K CHMOS DRAMs not only increase system density and simplify system design, but also improve system performance.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 713 *Average 714* *Low 715*

HERE IS ALL YOU NEED TO:



- Immediately interface any I/O channel compatible with IBM® Systems*
- Detect glitches, edge or error on tag line
- Monitor simultaneously up to 41 error or compare conditions
- Verify channel and device communications

COMCHEK INTERNATIONAL, INC.
8711 Pinnacle Peak Road, Suite 201-F
Scottsdale, Arizona 85255 Phone: (602) 585-3668

*PAT. PENDING

CIRCLE 73

*IBM® Interface cables included in the portable case lid.

©1984 AT&T Technologies, Inc.



UNIX™ System V from AT&T can help you close a sale, as well as open up new markets.

It's another reason why good business decisions are based on UNIX System V from AT&T.

More and more of your customers will be demanding multi-user and multi-tasking systems. The benefits of UNIX System V—portability, reliability, flexibility—make it the ideal software system for this lucrative market.

In addition, AT&T will support UNIX System V with the service, training, and documentation resellers need to prosper.

More choices, more opportunities

UNIX System V is virtually hardware independent. So you'll be able to configure systems using equipment from a variety of vendors.

You can use the flexibility of UNIX System V as a real selling point to customers who want to add hardware that is compatible with the machines they already have.

And over the long term, UNIX System V offers continuing opportunities to upgrade and expand your customers' systems. End-user flexibility translates into market opportunity for the reseller.

Software that sells

UNIX System V offers



another real sales opportunity in the area of customized software.

You'll be able to sell software for a wider range of configurations. You'll have more software to offer, too. Your programmers can concentrate on developing new packages instead of rewriting old ones.

For more information, send in the coupon for a copy of our free booklet, "Why Good Business Decisions are Based on UNIX System V."

UNIX System V. From AT&T. From now on, consider it standard.

AT&T

Please send the free booklet "Why Good Business Decisions are Based on UNIX™ System V."
Mail to: AT&T, P.O. Box 967,
Madison Square Station,
New York, NY 10159

Name _____

Title _____

Company _____

Address _____

City _____ State _____ Zip _____

My business category (check one):

☐ OEM/VAR ☐ Software House

☐ MIS/DP Mgt. ☐ General Business Mgt.

Other _____

UNIX System Licensee ☐ Yes ☐ No ☐ Don't know

CD0800CA

EHVQ™ SWITCHING POWER SUPPLIES WITH OPTIONAL BATTERY BACKUP



POWER SUPPLIES DESIGNED AS INTELLIGENTLY AS YOUR COMPUTER

Our EHVQ™ line of high quality compact switchers offer new flexibility in the design of fault-tolerant systems. In addition to many configurations, offering one to four outputs, the EHVQ™ family is available with battery backup to insure reliable DC power during brownouts or blackouts. Optional battery backup eliminates the size/cost penalty of traditional backup systems. Battery backup is field upgradeable. It plugs into your system's EHVQ™ power supply, as easily as adding more memory.

TTL compatible status signals allow the EHVQ™ to communicate with your system to prevent loss of RAM and disk data. EHVQ™ Switchers....

HELPING YOU BEFORE YOUR CHIPS ARE DOWN.™

EHVQ POWER SUPPLY

- 170 watt and 220 watt models
- Single, triple and quad output models
- Remote on/off
- Industry-standard chassis
- Built-in line filter

BATTERY BACKUP OPTION

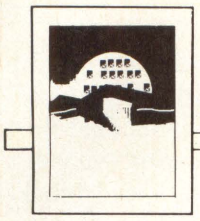
- Field upgradeable (plug in PC board)
- TTL status signals
- Jumper selectible shutdown modes
- On-board, float-type battery charger

TII® ELECTRONICS, INC.

226 Terminal Road, E. Setauket, New York 11733 516-751-6285

Circle 74 for literature Circle 75 to have applications engineer call

© 1983 TII ELECTRONICS



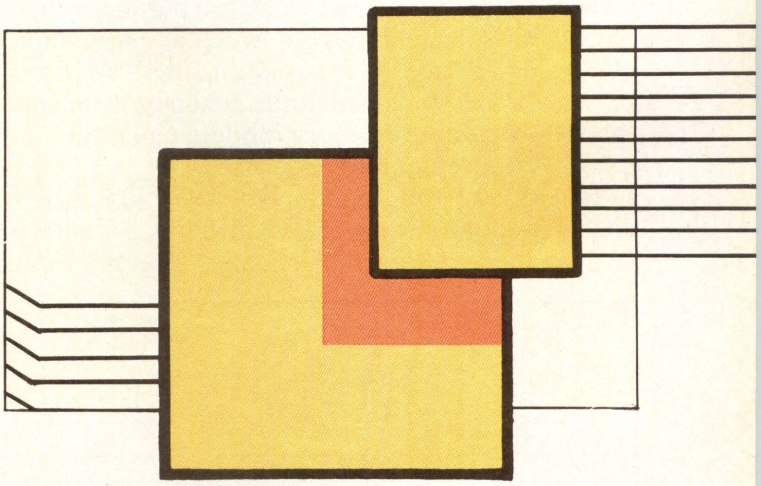
DUAL-PORT STATIC RAMs CAN REMEDY CONTENTION PROBLEMS

System bandwidth can be used up by contention problems. A true dual-port RAM can help by allowing simultaneous reads and writes by different processors.

by Michael J. Drumm, James B. Harris, and Michel Ebertin

There are many advantages when multiple CPUs and intelligent peripheral processors share the same bus. However, to fully gain these advantages, they must run at their top speed. This means being able to access the necessary memory locations immediately.

All arbitration schemes, however refined, involve waiting. With the use of a truly dual-port RAM, however, (eg, Synertek's SY2131), different processors can access the same memory locations simultaneously. Contention is handled at the byte level with onchip arbitration logic. Thus, the chances of any processor having to wait for access are greatly reduced.



To make the dual-port implementation possible, Argonne Systems, Inc has developed the multiple coupled processor (MCP) series of computers. This series is based on an architecture using a system of closely coupled processors residing in the same backplane and sharing common peripherals. The Tarch architecture enables different operating systems to run on different CPUs simultaneously, while sharing peripheral resources, and communicating at the application software level via a common bus.

Michael J. Drumm is currently product marketing manager at Argonne Systems, Inc, 525 Julie Rivers Rd, Sugarland, TX 77478, where he is responsible for product planning. Mr Drumm holds a BS in electrical engineering from Arizona State University.

James B. Harris is a development engineer at Argonne Systems, Inc, where he is responsible for storage control processor design. He holds an associate's degree in electrical engineering from Wharton College.

Michel Ebertin is vice president of engineering at Argonne Systems, Inc, where he is responsible for product design. He holds a BS in electrical engineering from Polytechnic Institute of New York and an MS in electrical engineering from California Institute of Technology.

The basic building blocks of Tarch are a CPU module containing a standard microprocessor (with different microprocessors requiring different CPU module implementations); a cached, dual-ported, memory module with up to 4 Mbytes of main memory and 16 Kbytes of cache; a storage control processor (SCP) module providing interface, buffers, control, and timing between the system's main bus and the mass memories' Small Computer System Interface (SCSI) bus; and an I/O processor module. These building blocks can be assembled singly, or in multiples, to create a system. Fig 1 illustrates a dual-CPU option.

The first CPU module offered by the company is designed around the Digital Equipment Corp J-11 microprocessor. This device, a single-package implementation of a PDP-11/70, executes the full PDP-11 instruction set. Used within the Tarch environment, this module executes existing application software faster than any other PDP-11-based system. Two such modules in a system (with their respective memory modules), give multiple users the ability to run simultaneously under two different operating systems.

As a complete 32-bit architecture, Tarch offers up to 4 Gbytes of physical memory space and features multiple buses. The first, the parallel system bus (PSB), is the main bus to which all basic building blocks are tied. It is based on Intel's Multibus II. Other buses include multiple auxiliary PHASTBUS (one for each CPU/memory module pair), the DEC Q-bus, and a unique I/O channel.

Two of the four building blocks, with widely differing data transmission bandwidths, simultaneously handle data. In particular, the SCP must contend on one side with a PSB bandwidth of 13.6 Mbytes/s and, on the other side, with a SCSI bandwidth of 1.5 Mbytes/s. Here, a static dual-port RAM plays a significant role in preserving the PSB's high bandwidth. Caused by its high speed block-mode transfers, the dual-port RAM buffers slow speed device data on the peripheral controller side.

Looking at the SCP operation

The SCP is an intelligent module providing the timing, control, and protocol conversions between the PSB and SCSI buses. It emulates the digital storage architecture and a TSV05 tape controller, while using the mass storage control process protocol.

The SCP handles all disk and tape I/O traffic for the MCP series of computers. Multiple SCPs can coexist in the same backplane, with each SCP handling multiple CPUs, and operating as though it had a dedicated I/O port and peripheral. The SCP handles all transfers to and from the SCSI bus while doing all mass storage control process and TSV05 protocol emulations, servicing the communication registers, processing interrupt messages to other PSB arbiters, and communicating to a CPU via direct PSB address space. National Semiconductor's NS32016 microprocessor, Advanced Micro Devices' AMD9516 DMA controller, NCR's NCR5385 SCSI controller, and six Synertek 2131 dual-port RAMs were selected to perform the SCP's tasks.

In the SCP's primary function, the PSB's high bandwidth must be coupled with that of the much slower SCSI bus. Thus, the buffer memory of each SCP provides simultaneous, unrestricted access to each bus, except where the same memory location is being addressed. Mass storage peripheral data

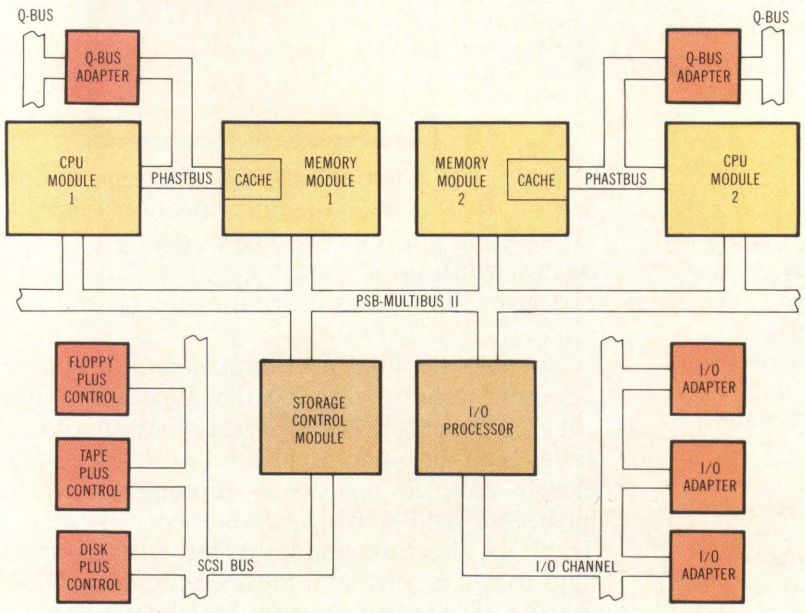


Fig 1 With the Tarch architecture, using two CPUs, all four types of modules exchange data via the parallel system bus (PSB). It is imperative to minimize contention on this bus and use its full potential bandwidth.

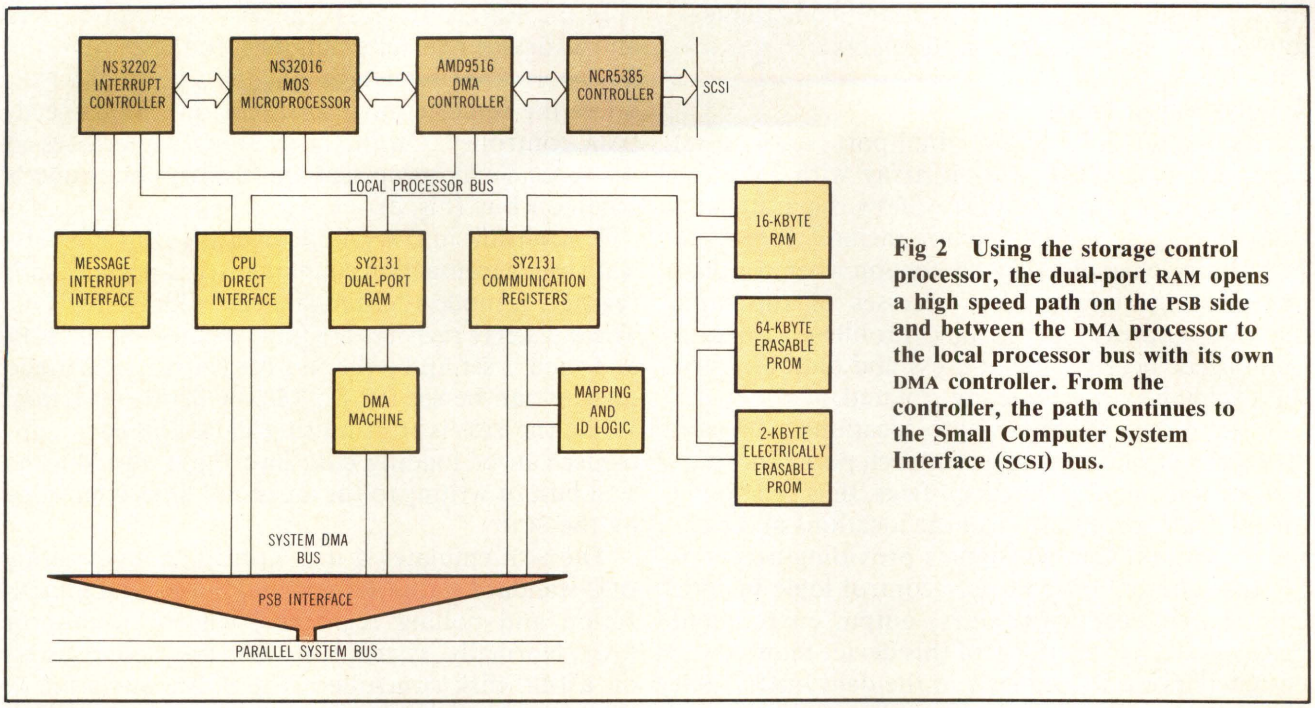


Fig 2 Using the storage control processor, the dual-port RAM opens a high speed path on the PSB side and between the DMA processor to the local processor bus with its own DMA controller. From the controller, the path continues to the Small Computer System Interface (SCSI) bus.

must be located, retrieved, and formatted into packets before transmission; the reverse is required for incoming data. Since the SCP is accessible to all CPUs in the system, it distinguishes data and allocates storage locations on shared peripherals, on the basis of the storage allocation parameters specified during system startup.

Testing determines that even a bit-slice microprocessor could not keep up with the data transmission rates while providing the necessary computational services of protocol formatting, request queuing, and managing data to/from the buses. The read/write cycles necessary to move data between the buses would be too slow.

System solution

The solution to this problem requires a DMA controller to handle data on and off the SCSI bus. This frees the microprocessor for protocol emulation. Also needed is a custom DMA controller on the PSB side to arbitrate for the PSB, and transfer data at the required speeds. A truly dual-port RAM decouples the two buses and allows simultaneous, asynchronous data transfers.

First, a DMA controller is added to the NS32016 microprocessor architecture to handle DMA transfers to and from the SCSI controller. This opens the bandwidth of the microprocessor by removing direct data transfer overhead between the SCSI channel and local buffer memory. The data transfer rate is limited only by the SCSI standard. If the NS32016 microprocessor alone were used to control data flow to and from the SCSI controller, a minimum of four instructions would be required per transfer. Data transfer to and from the SCSI bus requires minimal direct interaction between the microprocessor and the DMA controller. The NS32016 microprocessor places instructions and addresses for the controller in the chaining table of the local RAM. Then, the microprocessor issues a DMA transfer command to the SCSI controller, and the DMA controller executes the transfer.

A traditional double-buffer implementation, in RAM or with first in, first out (FIFO) buffers, would allow the DMA controller to load the buffer while data was being transferred onto the PSB. But available DMA controllers are too slow to service the PSB, and unsynchronized transmissions require decoupling the two processes. Therefore, a discrete-logic DMA controller is placed on the PSB side of the buffering. Note that FIFOs could provide the required decoupling, but their high access times make them too slow for the PBS speed requirements. The dual-port RAM decouples data transfers from the SCSI channel to the high PSB bandwidth. The SCP's PSB DMA controller allows data transfers between the dual-port buffer and the PSB-resident source/destinations.

This approach dictates the use of a high speed RAM buffer able to be accessed simultaneously from two sides. Lacking such simultaneous access, the microprocessor's ability to do emulation functions would be severely impaired. Also, potential benefits of the high speed PSB would be reduced significantly when the two communication processes need to be coupled. A conventional, dual-ported RAM, with switchable but exclusive access states (ie, when one port has access the other does not), would not be sufficient. The SY2131 dual-port RAM not only provides the necessary functional solution, it also

fits within the PC board "real estate" constraints—impractical with discrete components.

A truly dual-port RAM

The SY2131 is a truly, dual-port static RAM organized as a 1024- x 8-bit array with a 100-ns access time. Two I/O ports allow separate and simultaneous access to common memory locations. Internal logic handles data contention when conflicts arise during simultaneous accesses of the same memory address. Cross talk problems are also minimized, because the address and data buses of the two ports are physically isolated.

Writing data internally into locations 3FFH and 3FEH creates interrupt flags for each port. This triggers the left and right interrupt flags, thus prompting the microprocessors to read the locations and clear the interrupts. Control signals providing access to internal contention and port control logic include chip enable, read/write lines, output enable, and busy flags. The versatility of this device allowed the designers not only to use it in the data buffer, but also in device register emulator, and communication register applications.

In the SCP, four SY2131s are used as data buffers and as device register emulators. In these roles, the dual-port RAM interfaces on one side to the PSB, operating under control of a custom, high speed, DMA controller. On the other side, it interfaces to the NS32016 microprocessor and the DMA controller's common bus (Fig 2).

To transmit in DMA mode over the PSB, SCP software uses a transaction count register, a packet count register, a mode register, and an address register. These discrete registers are implemented on the SCP, and each is set up by the NS32016 (Fig 3). The mode register can be set for 8-, 16-, or 32-bit transfers. While the SCP is transmitting data, the DMA controller can be loading data into the RAM from the SCSI bus, or writing to the tapes and disks controlled by the SCP.

The SCP emulates a disk controller by creating DEC-standard status address and purge, and initialization and polling registers inside the dual-port RAM. Normally, system software sees these registers on a DEC disk controller. The SCP maps a set of these registers into the dual-port RAM for each CPU accessing it. Through the RAM's mapping, the SCP

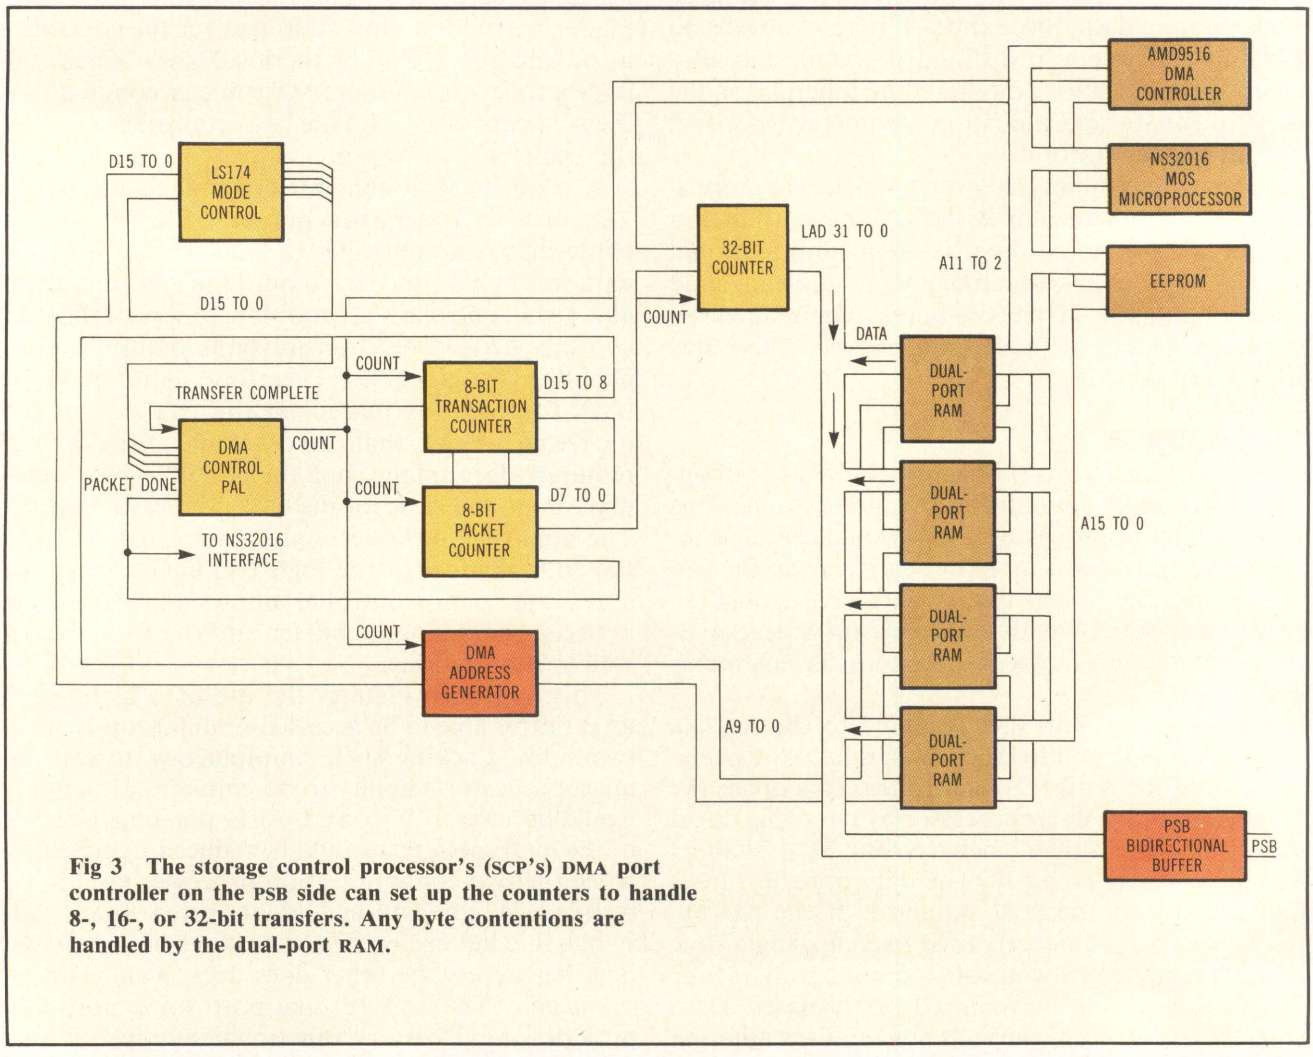


Fig 3 The storage control processor's (SCP's) DMA port controller on the PSB side can set up the counters to handle 8-, 16-, or 32-bit transfers. Any byte contentions are handled by the dual-port RAM.

How to debug in your favorite language.



Now you can debug in Pascal, C or Assembly. With Emulogic's new SLICE Symbolic Debugger.

First, sit back. Second, choose your language—C, Pascal or Assembly. Then type it in. And you're ready to start debugging.

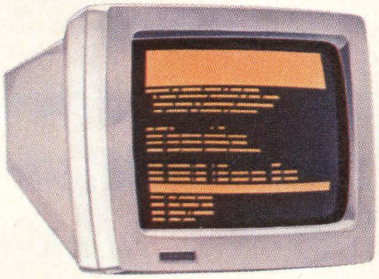
At least, that's the way you do it with SLICE™ (Source Language In-Circuit Emulator). And that's not the only way SLICE makes your life easier.

It's also friendly, interactive and non-intrusive—which means you don't need a software monitor. You'll be able to debug your code without us getting in your way.

What's more, if you change your mind about which language is really your favorite, SLICE lets you switch back and forth. Instantly.



It's true high-level debugging.



We didn't compromise when we designed SLICE, and you don't compromise when you use it. You can single-step by source line or chip instruction—and, if you want, step over calls to functions. Display real-time traces in high-level and assembly-level code. And manipulate stack, heap, register and static variables of all data types, including structures and arrays.

We've even put in another exclusive feature, TRACEBACK, that lets you examine all active functions and arguments. And another feature that lets you track movement of fast variables between registers and stack.

It's a real-time in-circuit emulator.

SLICE is a true in-circuit emulator—not just a software simulator. So you don't lose any of the hardware and assembly-level debugging power of Emulogic's universal MDS's. You still have



eight real-time multifunction breakpoints. "Fall-through" trapping. And full-speed, full-range memory mapping.

But SLICE makes them even better. Now you can break at a source line and use addresses, data and expressions symbolically. And select a trace of machine cycles, instructions or source lines.

Our Pascal and C compilers are true cross-compilers. Which means they produce executable code that can run on an Emulogic MDS or your target chip. Both compilers, of course, are designed to be used with SLICE. But, if you prefer to use a third-party compiler, we can provide you with all the information you need to make it work with SLICE.

SLICE packages are priced from $5,000 to $12,000, depending on your operating system and the compilers you select. They're now available for the 68000 and soon will be for the 68008, 68010, 8086, 80186, 80188 and 8088.



SLICE. Isn't it about time you started debugging in the language you're actually using?

For complete information, contact Emulogic,® Three Technology Way, Norwood, MA 02062, 800-435-5001 or 617-329-1031 in Massachusetts.

SLICE is a trademark, and Emulogic is a registered trademark, of Emulogic, Inc.

EMULOGIC®
MICROPROCESSOR
DEVELOPMENT SYSTEMS

European Distributors: Austria: Walter Rekirsch, (43 222) 235555; Denmark: Instrutek, (45 5) 611100; France: YREL, (33 3) 9568142; Sweden: Aktiv Elektronik AB, (46 8) 7390045; Switzerland: Instrumatic AG, (41 1) 7241410; United Kingdom: MSS, (44 494) 41661; West Germany: Instrumatic Electronic GmbH, (49 89) 852063.

CIRCLE 76

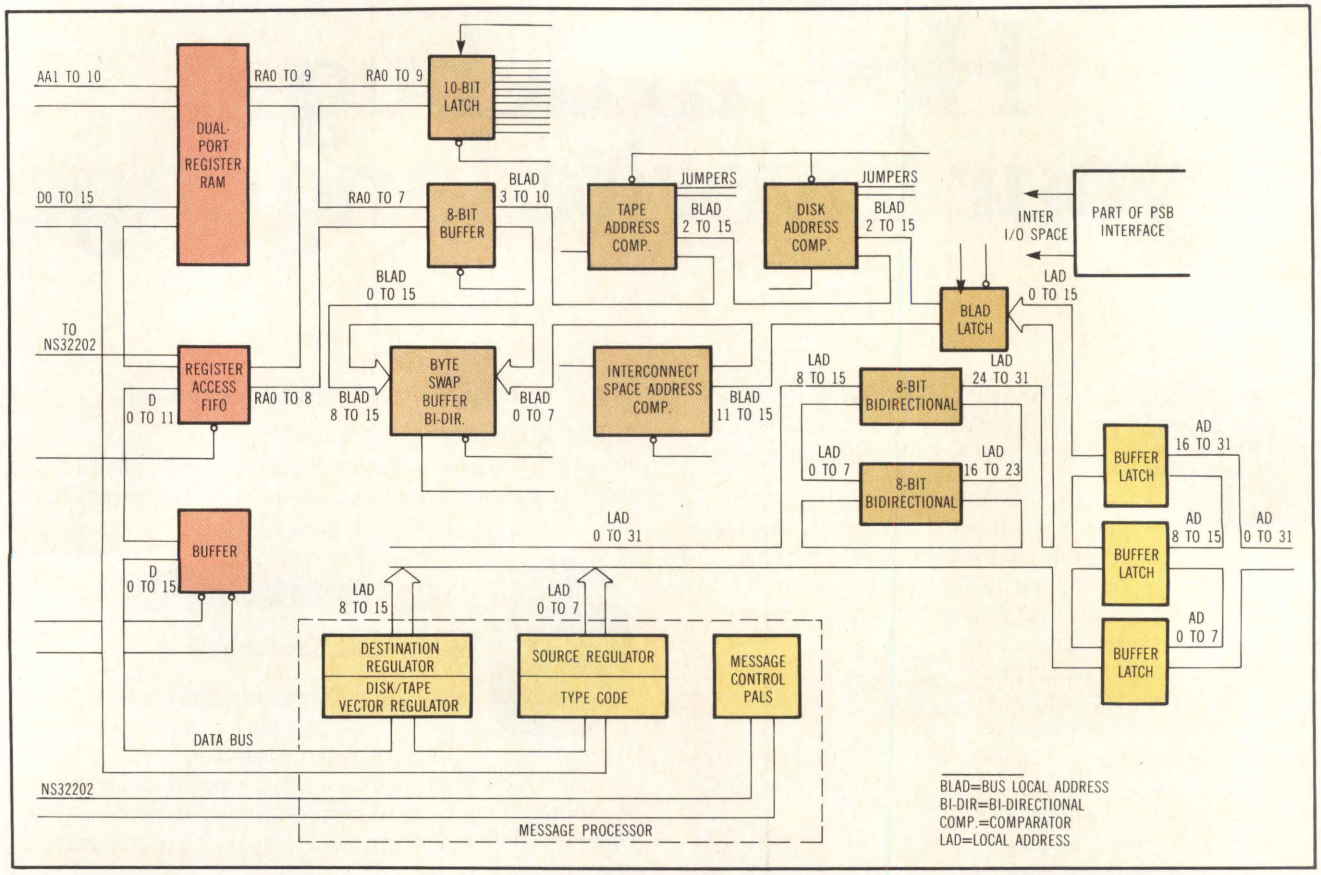


Fig 4 In this example of an SCP message processor and disk/tape registers, byte swapping for correct alignment occurs on the local bus. It is placed in easily accessed locations in the dual-port RAM during the transfer operation.

appears to the system as an n-ported controller with each port dedicated to a particular CPU and assigned an address in that CPU's I/O space, where N equals the number of CPUs in the system.

Communication registers

Two SY2131s dual-port RAMs serve as communication registers. Information stored in these registers identifies the peripheral operation (tape or disk) and the requesting or receiving port in a data transfer operation. One side of the dual-port RAM communication register array (1-K x 16 bits) is interfaced to the NS32016 address and data bus (Fig 3). A busy signal output from the SY2131 generates wait states in the NS32016 or the DMA controller cycle if simultaneous writes to the same byte location occur. A 10-bit latch captures information to address the dual-port RAM, based on the owner ID number that tells which module it is and whether it is tape, disk, or interconnect access. The Mutltibus II protocol also provides interconnect access during system startup to aid in system configuration management.

Dual-port RAM space is partitioned to separate interconnect space operations from I/O space operations. Within the I/O segment, additional partitioning separates disk and tape operations, with sets of 16-bit registers used in each category for read/write functions. These registers emulate reserved addresses in the DEC I/O space. There are complete sets of these registers reserved for each of the 20 CPU modules that can reside on the PSB. Using this scheme, each CPU in the system has its own set of communicaton registers implemented within the DEC I/O space for tape, disk, and read/write operations. The owner ID is captured for the interconnect space (512 x 8 bits), which is also mapped in the dual-port RAM, using Multibus II specifications. When accesses are made in the interconnect space, this register set keeps track of the originator or destination ID.

A register FIFO acts as a buffer for multiple, high speed requests to access the communication registers, For each data transfer request, the FIFO holds the owner ID, a read/write flag, and a flag that selects interconnect space or I/O space. As soon as the FIFO stores any information, an interrupt issued to the microprocessor notifies it to read the FIFO until it is empty. This action allows multiple register accesses to occur before the microprocessor has a chance to service the transfers.

To drive out to the PSB from the dual-port RAM, the microprocessor generates the required 32-bit PSB address. The SCP tags on the arbitration ID of the receiver, and knows where in memory to store the data/command end messages from the mass storage control process or TSV05 protocol.

Activating the DMA port initiates the PSB's request phase. The command line state is changed to indicate the reply phase, and data transfer begins. At this point, output enables and chip selects are generated to the dual-port RAMs. The interface is designed to do 8-, 16-, or 32-bit transfers. Byte swapping on the local data bus correctly aligns words and bytes transmitted on odd boundaries. The swapping registers involved are transparent to the software (Fig 4). A 16-bit command and address register holds the physical, dual-port RAM, address, and command information. Bits 0 to 9 contain the DMA RAM address.

Maintaining packet and bus transaction counts

During DMA data transfers, a packet count and a bus transaction count are maintained. Eight-bit counters are used, for a maximum count of 256 packets, with 256 transactions in each packet (maximum of 256 Kbytes data transferred with one initialization). The packet counter is loaded with the predetermined number of data packets to be transferred. For each PSB access, one packet is sent and the counter decrements by one. When the packet counter equals zero, a completion interrupt is generated to microprocessor and an "end of cycle" is sent on the PSB.

The transaction counter is also loaded with a predetermined number representing the amount of bus transfers allowed. The counter decrements by one for each executed transfer. A count related to long word transactions (32 bits) is also kept. This count ensures properly continued data transfer when multiple packets are sent to the PSB. A RAM address counter, coupled with the 32-bit dual-port RAM address register, counts the long word transactions.

With the SY2131 dual-port RAM, the company produces a high performance SCP that can service multiple CPUs using industry-standard buses with widely varying speeds. A discrete logic implementation would have involved five to seven times the chip count, and still would not have achieved full simultaneous access from both ports. Such a discrete implementation to achieve decoupling (a FIFO) would have resulted in a design that could not satisfy the PSB speed requirement. The performance and versatility of the RAM is encouraging in a number of further design situations, including a cache memory for a magnetic tape controller.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 716 *Average 717* *Low 718*

SYSTEM V

TRAINING

For 15 years, we've taught our own people to use the UNIX™ System. Now we can teach yours.

WHY AT&T FOR UNIX SYSTEM TRAINING?

AT&T offers the most current and comprehensive training on UNIX Systems.

AT&T provides the best learning environment; one terminal per student; evening access to facilities; and expert instructors.

AT&T has the breadth of courses your staff needs to unlock the full power of UNIX System V.

AT&T courses signal your commitment to improving productivity with high-quality training for your employees.

AT&T COURSES OFFER:

The same training and methods we use to teach the UNIX System to our own people.

Rigorous classes designed to teach specific skills for job-specific applications.

Five areas of instruction ranging from introductory to advanced levels for Managers/Supervisors, Users, Systems Administrators, Applications Developers, and Systems Programmers.

Frequent class offerings so you won't have to wait for the courses you want.

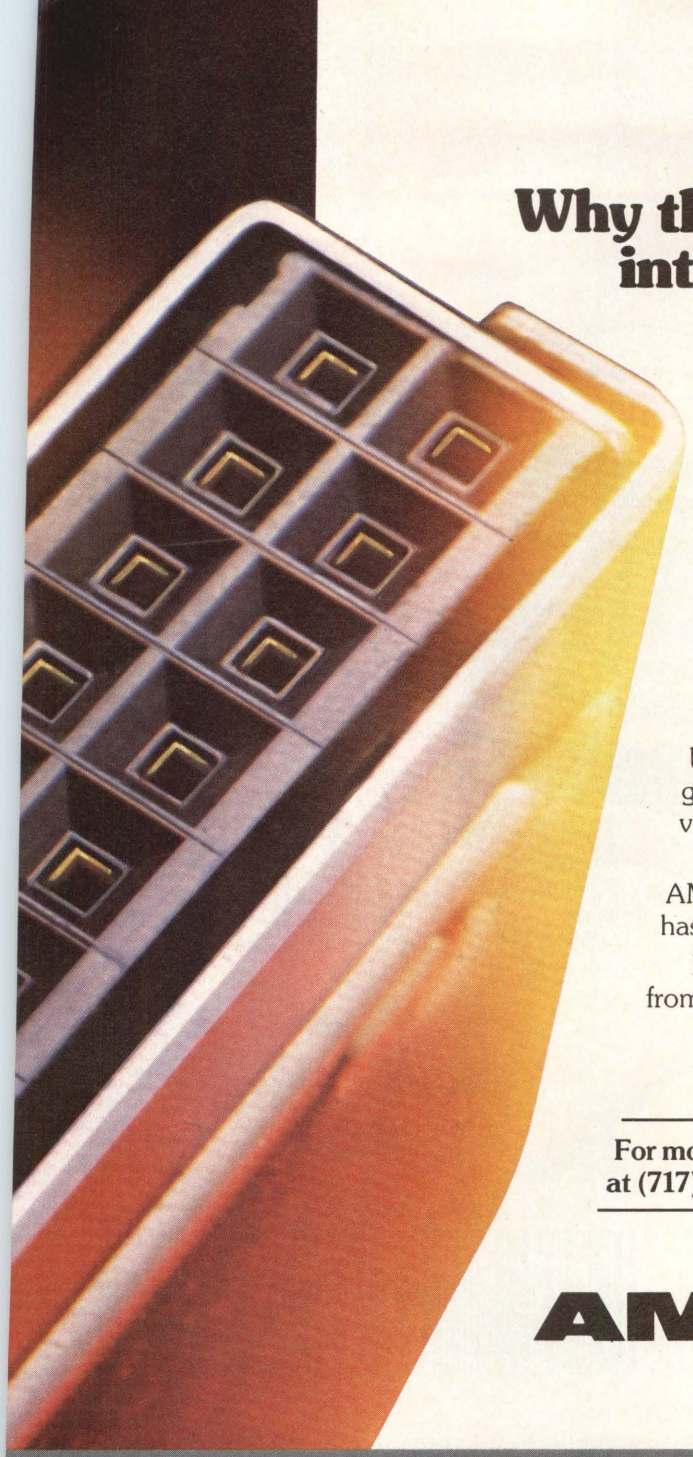
Conveniently located training centers in Princeton, NJ; Columbus, OH; Lisle, IL; and Sunnyvale, CA. Or we'll bring our courses to your company and hold the training at your convenience.

For more information, a catalogue, or to register for classes, call **1-800-221-1647, Ext. 23.**

AT&T

©1984 AT&T Technologies, Inc.

CIRCLE 77



Why the most popular interconnect system in the world got that way.

People choose the AMPMODU modular connector system for its tremendous flexibility.

Modular design and construction give you the versatility you need to connect just about anything to anything. Board-to-board, wire-to-board, wire-to-wire.

Example: 4000 part numbers on headers alone.

People also choose AMPMODU connectors for their big list of options and value-added features. Accu-plate gold plating. Solderless compliant pins. Mass-termination versions, too, with shielded designs for EMC.

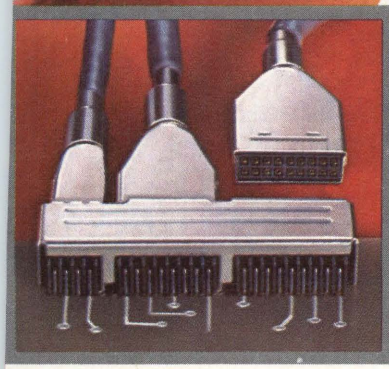
And everyone who chooses AMPMODU connectors gets AMP quality, AMP reliability. Our dual-cantilever beam design has proven itself dependable in over sixteen years of use.

Proven technology. Proven flexibility. Proven savings, from AMP.

For more information, contact the AMPMODU Connector Desk at (717) 780-4400. AMP Incorporated, Harrisburg, PA 17105.

AMP means productivity.

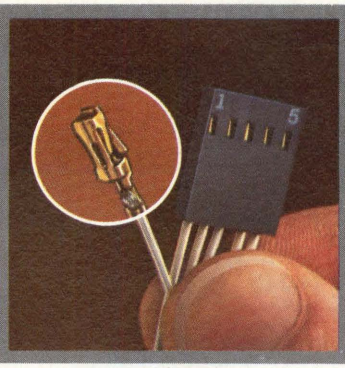
CIRCLE NUMBER 10



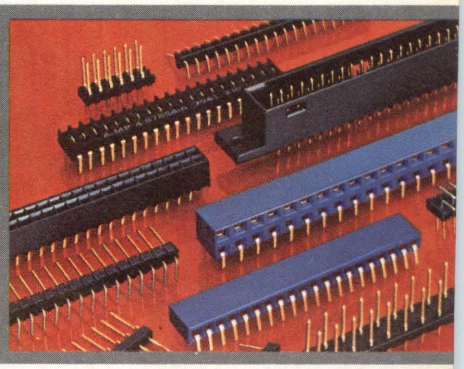
Shielded mass termination versions available for EMI/RFI compliance.



AMPMODU MT connectors use reliable insulation-displacement technology, to mass terminate discrete wire, twisted pair, jacketed or ribbon cable.



Crimp-snap technology for complete flexibility in wire-to-board, wire-to-wire applications.



Standard, shrouded, and breakaway headers for board-to-board mating. Compliant-pin versions for solderless insertion.

AMP and AMPMODU are trademarks of AMP Incorporated.

Break the Pencil Input Bottleneck...

With Low Cost Bar Code Terminals





EAN



2 of 5 (interleaved)



Code 39



CODABAR



2 of 5 (standard)

Fast, accurate, "direct-to-CPU" data collection at every work station is practical now! Low cost, operator friendly, Microterminal™ terminals give you display plus keyboard and bar code inputs in a tough 8-1/2" x 4-1/2" x 1-1/4" waterproof package that fits almost anywhere. You can set the reader for any of five standard bar codes with a rear panel switch. Five 50-character buffers handle bar code inputs. Using industry standard interfaces, up to 63 Microterminal terminals can work on a single serial line to interface with most computers.

TM71B - alphanumeric entry keyboard - and **TM77B** - numeric entry keyboard - are expanded versions of industry tested/accepted TM71 and TM77 Microterminal models. With intergral bar code reader decoder and wand, units are priced at $950 in 100's!

Request full information on these feature-loaded 100% solid state terminals that can replace bulky CRT's and printers or work where these fragile devices can't.

Data Acquisition & Control Systems Division
3631 E. 44th Street, Tucson, AZ 85713 (602) 747-0711

BURR-BROWN®

Putting Technology To Work For You

(205) 882-0316, (206) 455-2611, (213) 991-8544, (214) 681-5781, (215) 657-5600, (216) 729-3588, (301) 628-1111, (301) 251-8990, (303) 663-4440, (305) 365-3283, (305) 395-6108, (312) 832-6520, (313) 474-6533, (314) 291-1101, (315) 699-2671, (315) 853-6438, (316) 942-9840, (317) 636-4153, (319) 393-0231, (404) 447-6992, (408) 559-8600, (412) 487-8777, (505) 883-3668, (602) 746-1111, (607) 785-3191, (612) 884-8291, (614) 764-9764, (617) 444-9020, (713) 988-6546, (714) 835-0712, (716) 544-7017, (716) 889-1429, (801) 467-2401, (805) 496-7581, (813) 885-7658, (913) 342-1211, (914) 964-5252, (919) 722-9445, **CANADA**: (403) 230-1341, (416) 678-1500, (514) 731-8564, (613) 722-7682

CIRCLE 78

HIGH ENDURANCE EEROMs OPEN SYSTEM OPPORTUNITIES

Because they can endure more read/write cycles, the latest electrically erasable ROMs can handle a wider variety of applications.

by Charles Furnweger

Few memory devices have evolved as rapidly as the electrically erasable ROM. Since its introduction in the late 1970s, this evolution in device technology, memory cell design, and device architecture has given the system designer an alternative to nonvolatile semiconductor memory devices.

The major advantage of the EEROM over other nonvolatile memory types is its ability to be programmed directly in a user's system. In-site programming eliminates removing a device from a system, erasing it with an ultraviolet light, reprogramming it in a separate programmer, and inserting it back into the system. Generating complicated waveforms and software algorithms to program the device is also unnecessary.

Early EEROMs, used primarily in specialized applications, did not find wide acceptance. Competing PROM devices were acceptable system alternatives. But, EEROMs have improved and no longer require high voltage inputs and sculptured programming waveforms. They work off a single 5-V power supply and are fully compatible with TTL levels. In addition, density, write times, access times, and endurance—measured by the number of read/write cycles devices can withstand without sacrificing reliability—have continued to improve (Table 1).

Charles Furnweger is a senior applications engineer at Seeq Technology, 1849 Fortune Dr, San Jose, CA 95131, where he is responsible for nonvolatile memory and communication products. Mr Furnweger holds a BS in electrical engineering from the University of Florida.

TABLE 1
EEROM Evolution

Parameter (units)	1980	1982	1983	1984
Erase/Write Time (ms)	50	10	1	.5
Density (Kbits)	16	16	64	64
Endurance ($\times 10^6$ cycles)	0.01	0.01	0.01	1

Note that EEROM devices are not always replacements for other nonvolatile semiconductor memories. Rather, they constitute a distinct product type that offers a unique set of capabilities to the system designer. These capabilities result from advances in circuit design and fabrication, Q-cell memory cell circuit design, and device architecture. The first two are reflected in EEROM performance. Device architecture simplifies EEROM use in a system. The increased endurance of the latest EEROMs, such as the million-cycle life of the Seeq 5516A, is

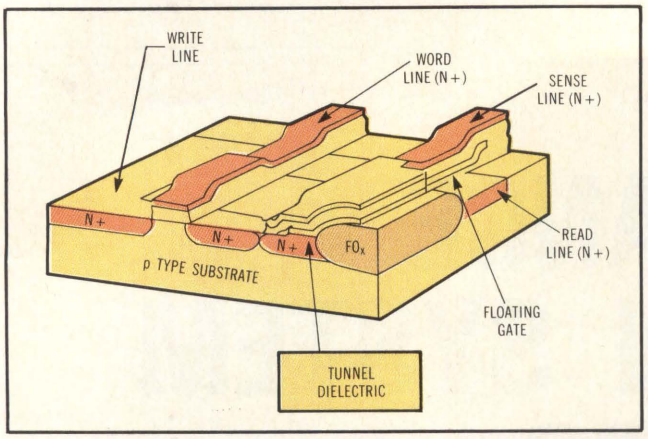


Fig 1 Advanced fabrication technology (Oxinitride) and a memory cell circuit design technique (Q-cell) increase EEROM speeds and endurance.

a result of both device technology and memory cell development.

One advancement in EEROM technology is Seeq's Oxinitride fabrication process (Fig 1). This process allows very thin, highly reliable oxides to be grown for the writing portion of the memory cell. As a result, the current available to write a charge onto a floating gate is significantly increased over that of other processes. This increased current reduces the time required to charge the cell to the appropriate level. Thus, the time required to reliably write to a particular location is also reduced.

Architecture and ease of use

EEROM architecture is one factor that determines ease of use. Features such as latches and timers built onto the chip, aim to benefit the system designer.

Latched devices store addresses, data, and control inputs available at the beginning of a write cycle. Timer devices automatically measure the elapsed write time, from 2 to 10 ms, depending on the device and manufacturer.

In general, it is easier to interface latched devices into systems that simultaneously require write operations for a large number of EEROMs. With latched devices, the user can time all devices from a common software time out, without waiting for each timer to notify the system of completion. Systems with limited board area and a few EEROMs are easier to design with timer-type devices. With these, the designer writes to the devices and waits until they respond.

TABLE 2
Memory Partition

Function	Memory Type	Major Features
Basic control code	ROM	Bootstrap, basic I/O system, firmware control
Application code	EPROM	Infrequent changes, high density, low cost
Application code patches, parametric data modification, configuration data	EEROM	Frequent changes, can be updated in system
Temporary data storage	RAM	High speed and density, low cost

The hardware and software required to interface EEROMs to a system depend on system bus and processor characteristics. The device's role is equally important.

In most systems, memory space is divided among a number of device types. Type choice depends on a number of factors, including the need for nonvolatility or modifiability. For a typical application, many different device types are required (see Table 2).

In sections where density and cost are important, ROM and EPROM devices are most appropriate. EPROM stores most of the application code because of the high probability of later changes. EEROMs also patch-in changeable portions of code.

EEROM used in the code-patching section of the memory makes use of its capabilities as a read-mostly, write-frequently device. EEROM provides for remote patching of application code, updating of parametric data, and downloading of control code. This design is even more valuable when a series of these systems is connected in a network (Fig 2). The interface may be via local area network, modem, or a combination. Control programs and data can be passed to nonvolatile memory without physical intervention, ie, removing and replacing memory devices.

Remote patching

A challenging problem, especially in process control applications, is designing the system for simple service and repair. Any changes to correct or improve program control code can cause significant problems. But with a little foresight in partitioning and designing the application code, changes can easily be made.

The basic requirement for code patching is modular code with pointers to segments in EEROM (Fig 3). By storing the pointers in EEROM, changing the pointer to show the location of the new code changes the code. Of course, the modified code itself is also in EEROM.

In a networked system, changes can be made to all the systems on the network from a controlling node. It is even possible to make these modifications to systems in the field, thus reducing the cost of field service. A similar technique can update data in a system. This is useful in process control and configuration applications, where changes in code and data are often required.

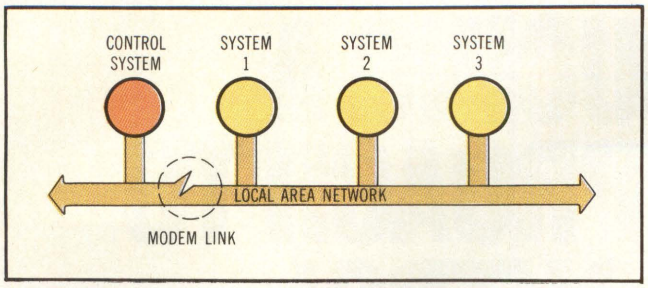


Fig 2 Electrical erasure benefits are brought out in a multisystem application, where code and data can be transmitted simultaneously over a local area network or modem.

These benefits can only accrue if the EEROM can endure a large number of erasure and writing cycles. After reaching its limit, the device must be replaced or the system can no longer maintain its performance.

Useful system life is the amount of write cycles the EEROM can endure. System reliability requirements and the frequency of write operations determine the endurance needed. Frequency of write operations is the number of times a particular EEROM location can be accessed for a write operation in a specified time period. System life can be estimated by dividing the endurance of the EEROM by the frequency of write operations.

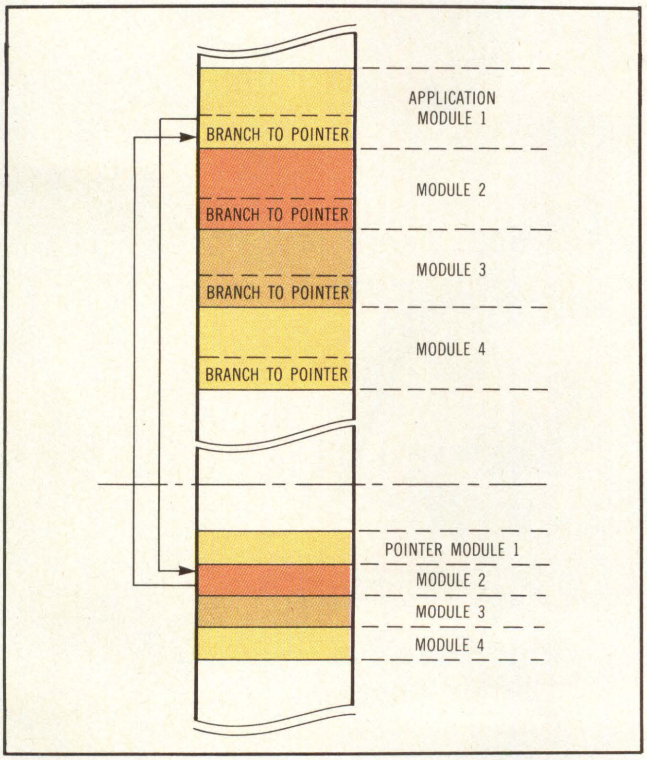


Fig 3 A simple way to change software remotely is to store code in modules, then change pointers and implemented code. Here, the end of application module 1 branches to a pointer register, which directs the processor to application module 2. Changing the pointer to application module 4 via modem link changes system operation.

In the example system, remote software patching, parametric data updating, and code downloading capabilities require different endurance specifications. Endurance requirements for patching application code are very small if the software writer is to have a long career. So this activity has little effect on the endurance requirements of the system as a whole. In process control applications, data might be updated as often as every half hour—48 times a day. New control code can be downloaded twice a week.

Because data is updated far more frequently than code, the data is placed in segments in EEROM. As the segments reach their endurance limits, new segments are defined and used. This procedure continues until no available segments remain.

Generally, the downloaded code is a significant portion of the EEROM's total available memory. As a result, it is not possible to increase endurance by segmenting memory. It is in this area that the benefits of high endurance EEROMs become apparent. At 100 writes per day, an endurance of 10,000 cycles translates into a system life of only three months. Increasing endurance to a million cycles increases system life to a year and two months.

System reliability

However, endurance should not be confused with reliability. Endurance measures how long a device will operate; reliability measures how well it operates. A device can have high endurance, but fail repeatedly (have low reliability) throughout its life.

Higher EEROM endurance improves overall system reliability. Many of today's systems do not need endurance greater than 10,000 cycles. In these systems, infrequent updating does not approach the 10,000-cycle limit, even over a long lifetime. But longer-life EEROMs may be needed for high reliability.

The expected failure rate for an EEROM can vary from 1 to 10 percent per 10,000 cycles, for devices with an endurance of 10,000 cycles. The higher endurance devices have a considerably lower failure rate (from 0.01 to 0.1 percent per 10,000 cycles). Even if a system does not require high endurance, it may require high reliability.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 719 *Average 720* *Low 721*

INTRODUCING A BETTER APPROACH TO ARMS CONTROL.

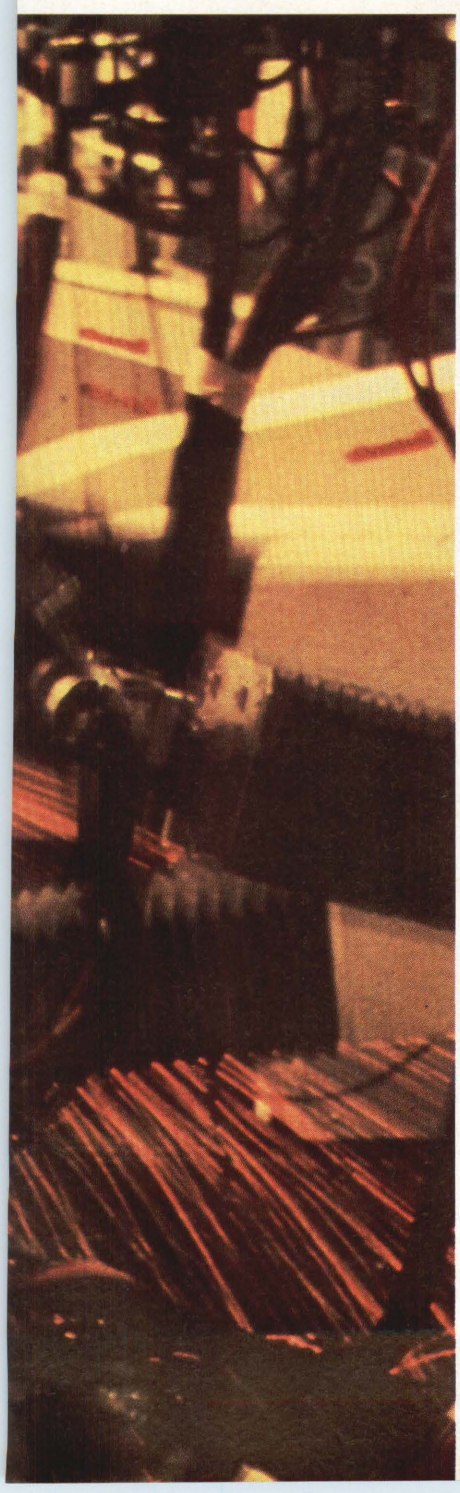


It's called the Intel approach.

Our 286/310 multitasking supermicro system. Teamed with our iRMX™86 real time operating system.

Together they provide OEMs in the industrial control and factory automation market with a system that's remarkably fast and flexible.

First, let's talk fast.



Our iRMX 86 operating system features priority driven interrupt management with ultra-fast context switching. And an event driven nucleus that manages multiple system tasks. All of which makes for lightning-fast real time response.

Then there's the 286/310 supermicro itself. Based on the most powerful high performance microprocessor in its class, our iAPX 286. Add to that the turbocharging talents of our 80287 floating point processor, which boosts numeric capabilities by 100x.

And add to all that our multiprocessing architecture. Which greatly enhances system speed by off-loading the CPU.

The result is scorching, minicomputer performance. Without a scorching, minicomputer price.

Now, let's consider flexibility.

Our iRMX 86 operating system is flexible because it's modular. So you only have to add those capabilities your customer's application demands.

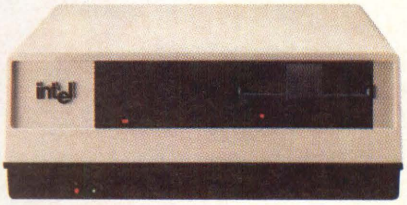
It also supports all of the popular languages (like FORTRAN, Pascal, PL/M, BASIC and C).

Making it flexible enough to accommodate the most diverse applications. A fact that over 2000 licensed RMX OEMs can attest to.

But the 286/310 supermicro is not only flexible, it's open. Giving you rapid access to the latest VLSI technology. Like our new BITBUS™ interconnect which provides low-cost, easy to implement distributed control in the factory.

And our open design allows easy integration of over 1200 different MULTIBUS™ boards available from over 200 suppliers. Letting you plug in I/O, data collection, control and communications capabilities.

So that's what the 286/310 real time supermicro system is all about.



Intel's 286/310 real time supermicro starts at under $10,000.

Speed. Flexibility. Simplicity. And the world's most advanced VLSI technology. Letting you easily combine chips, boards and systems — whatever you need. Plus complete maintenance. Even for the non-Intel parts of your system.

Call us toll-free for a copy of "How to Select a VLSI Operating System" and more information on the 286/310. (800) 538-1876. In California, (800) 672-1833. Or write Intel, Lit. Dept. #S 11, 3065 Bowers Ave., Santa Clara, CA 95051.

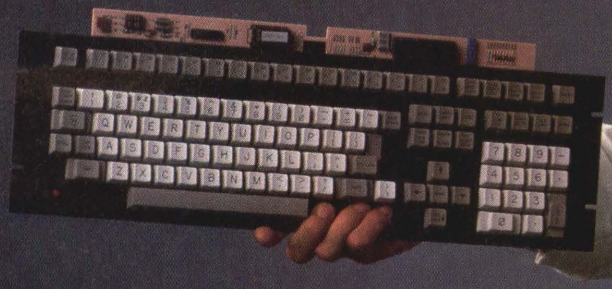
We'll be happy to give you a hand in winning the arms race.

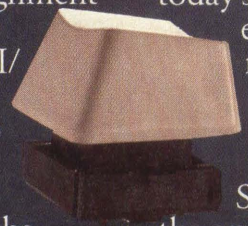
intel®

©1984 Intel Corporation

CIRCLE 79

You can put a better keyboard in your enclosure without modifying the enclosure.





Stackpole's new KS-600E keyboard–featuring snap-in modular switches–gives you a low-cost encoded keyboard that will retrofit into your present enclosure without retooling.

The discrete snap-in modular switches assemble into a metal plate to give you any kind of main or ancillary array layout you need. The metal plate assures alignment of the switches, adds rigidity and aids in RFI/EMI suppression.

You don't have to tool up for a new enclosure or modify the present one. We'll do the retrofitting for you on the backplate, and all you do is drop it in and connect it up.

And we can provide our custom encoded scheme, or design to your proprietary encoded format. Plus we let you choose between tactile or linear feel–at no additional cost.

The ergonomically designed full travel switches have a total stroke of .140 (3.6 mm nominal), are rated at 50 million cycles, are wave solderable and field repairable. The switches contain Stackpole's patented twin bifurcated contacts, with four make-point redundancy, which assures you of switch closure. And today's sophisticated electronics make mechanical switch operation reliable and dependable...at lower cost.

Stackpole also gives you the option of purchasing our modular switches individually, or in sets of configured arrays for your in-house assembly.

We have a variety of keycap styles. And we offer a major cost saving through our large legend library...tooled and in place...including the most popular international legends.

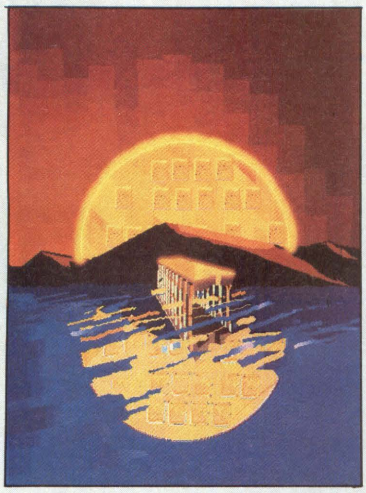
Our Application Engineers are ready to help, our domestic and off-shore facilities keep our prices attractive.

Call or write us today for complete information.

STACKPOLE®

Stackpole Components Company
P.O. Box 14466, Raleigh, NC 27620
Telephone: 919/828-6201

CIRCLE 80



Special report on minicomputer operating systems

163 Introduction

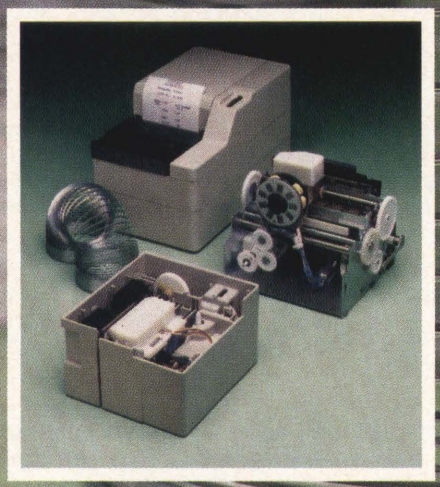
165 Operating systems for minis emphasize Unix compatibility
by Harvey J. Hindin—Realtime and fault tolerance capabilities are necessary features of any operating system that will meet 1985's minicomputer needs, particularly in applications such as robotics and process control.

177 Operating system features real time and fault tolerance
by Bob Snead, Frank Ho, and Bob Engram—A third Unix operating system layer, known as the kernel process, is the key software design innovation.

189 Fault-tolerant mini needs enhanced operating system
by Samuel D. Glazer—Unix System III application programs run in a user-transparent, expandable minicomputer environment.

205 Dual port solves compatibility problem
by Ross A. Bott—A single operating system incorporates both System V and 4.2 BSD Unix versions to solve program development/portability problems.

223 Minicomputer system offers time-sharing and realtime tasks
by Phil Sherrod and Stephen Brenner—An adaptive job scheduling algorithm, and both data and directory caching help optimize minicomputer operating system performance.



FLEXIBILITY

NCR 40-column printers...adaptable to any OEM application, by design. Your system shouldn't have to bend to accommodate a printer. The printer should conform to your system. NCR's economical, 40-column printers are the most accommodating mechanisms and standalone units in the world. Slip printing. Data logging. Receipt and journal printing. We've designed each of our printer configurations to be flexible, compact, rugged and easy to install. Today over one million units are in use around the world. To expand the limits of your design capabilities, call (800) 222-1235. In Ohio, (513) 445-2380.

NCR
OEM MARKETING DIVISION

NCR 40-column printers...designed for demanding applications.

CIRCLE 181



SPECIAL REPORT ON

MINICOMPUTER OPERATING SYSTEMS

Minicomputer technology is taking on more importance for computer system designers than ever before. Not only are minicomputers growing in power, they offer features not imagined just a few years ago. These features include, for example, realtime and fault-tolerant operation.

While the inexpensive microprocessor plays a major part in the new minicomputer functions, software—unique operating system software in particular—drives the hardware that is available almost as a commodity product. Indeed, minicomputer operating systems, like their microcomputer operating system counterparts (the subject of last month's Staff Report), are more accurately described as operating system environments.

These minicomputer operating system environments are far more than *ad hoc* software routines that turn a minicomputer on and off, run the I/O, and take care of resource allocation (the classic functions of operating systems). They are an integrated software resource, providing a rich set of services to computer system designers for a more cost-effective product. In addition to realtime and fault-tolerant capabilities, there are operating system software services to accommodate distributed hardware architectures, multiprocessing and multiprogramming, multitasking, time-sharing, networking, and program development.

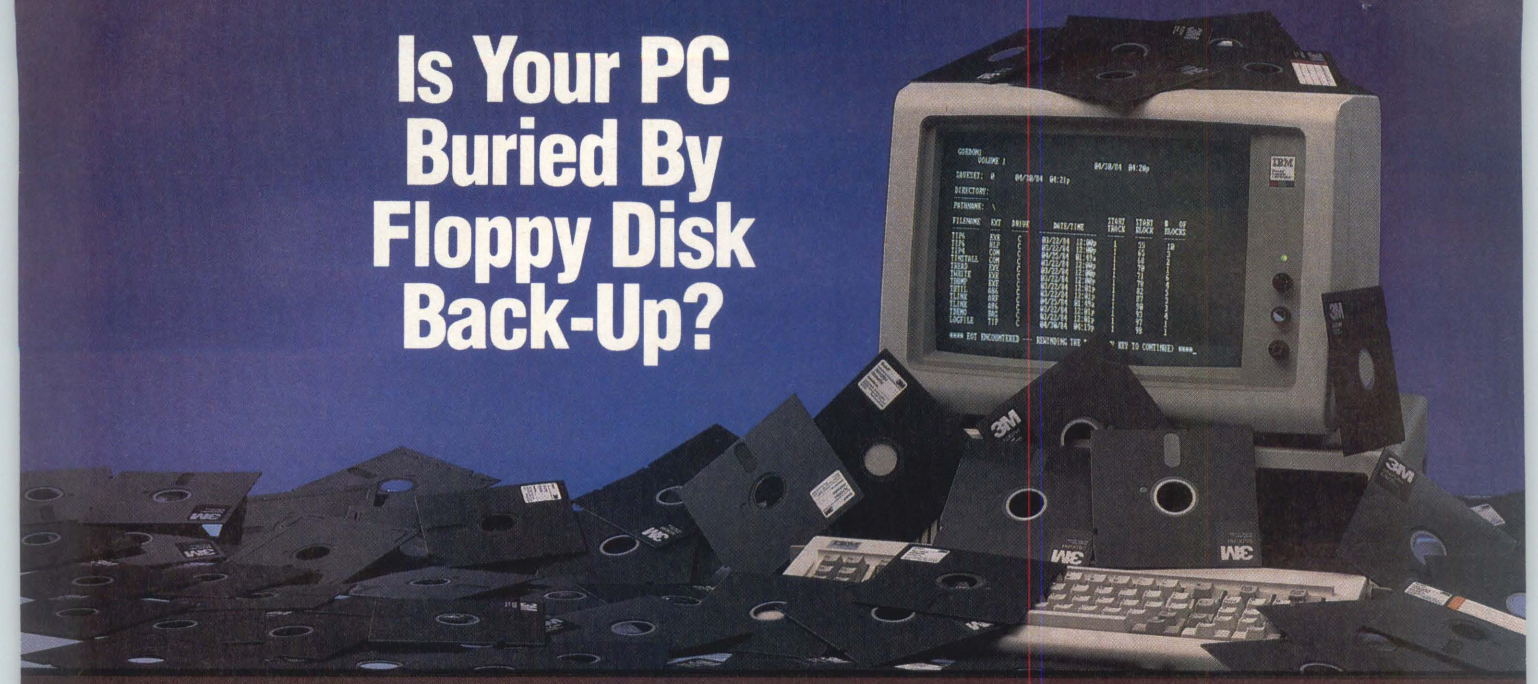
Each minicomputer manufacturer offers its own proprietary operating system. And, most often, some version of AT&T's Unix, whether licensed or not, is offered either by the minicomputer vendor or a third party. In fact, Unix has become the *de facto* standard minicomputer operating system—as much because there is just no other well-marketed contender that meets the needs of more than one vendor as because of its fine operating system attributes.

The contributed articles in this Special Report concentrate on what has been introduced this year to provide minicomputers with operating system environments. Accurately representing what is going on in the minicomputer industry, the articles are concerned primarily with combining Unix, real time, and fault tolerance in one state-of-the-art operating system.

Minicomputer system designers can learn from these articles what operating system software they will have as the basis for minicomputer system designs for the 1985 and 1986 marketplace. This evergrowing marketplace for the networked, integrated office and factory of the future demands an operating system environment that provides both user friendly functions and user friendly minicomputers.

Harvey J. Hindin
Special Features Editor

Is Your PC Buried By Floppy Disk Back-Up?





Dig Out With Model 70 PC Cartridge Tape System

Digi-Data's Model 70 PC cartridge tape system lets you back-up your PC's Winchester drive without getting buried in floppies. One cartridge holds 16.5 Mbytes of data, more than you can put on 51 floppy disks! And you can back-up your 10 Mbyte PC XT® drive in less than 15 minutes of *unattended* operation. That is a small fraction of the time it would take with the PC's floppy, with you standing there changing disks.

Our file-oriented structure makes it possible for you to back-up only what has to be changed on your disk. That saves still more time and storage capacity.

70 PC comes ready to run on your IBM PC XT and most compatibles like Compaq, Columbia Data Products and Eagle. Complete with controller board, driver software and cables.

Digi-Data's products also include 1/2 inch and 1/4 inch start/stop and streaming drives and systems.

CIRCLE 82

DIGI-DATA CORPORATION

®... First In Value

8580 Dorsey Run Road
Jessup, MD 20794
(301) 498-0200 Telex 87580

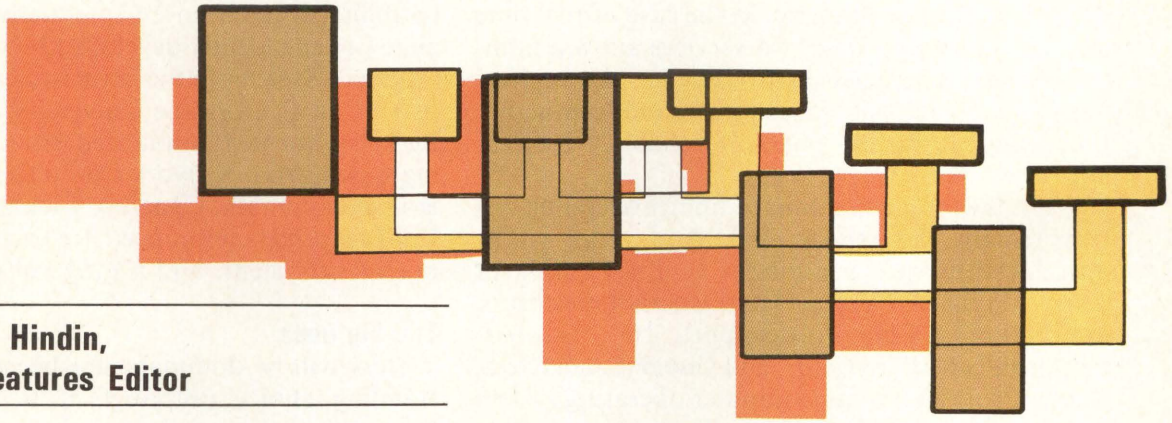
In Europe contact:

DIGI-DATA LTD
Unit 4
Kings Grove
Maidenhead, Berkshire
England SL6 4DP
Tel. 0628-29555-6
Telex 847720

® PC XT is a registered trademark of IBM Corporation

OPERATING SYSTEMS FOR MINIS EMPHASIZE UNIX COMPATIBILITY

Realtime and fault tolerance capabilities are necessary features of any operating system that will meet 1985's minicomputer needs, particularly in applications such as robotics and process control.



by Harvey J. Hindin,
Special Features Editor

For the second half of this decade, minicomputer operating systems will sail a clear course. They will emphasize Unix (and/or Unix look-alikes), realtime capabilities and facilities, and the ability to work with fault-tolerant hardware. Minicomputer systems for the computer system designer, be they standalone or networked, will run Unix or one of its clones as either a first-choice or alternative operating system. And, they will handle realtime applications as a first choice or as an option. The systems will allow varying degrees of fault tolerance. (These days, smart hardware is inexpensive enough to perform backup functions, even though it cannot be fully utilized for computing chores.)

Until now, Unix has comprised utility programs, application software, and development tools. But, with the addition of realtime and fault-tolerant software features, minicomputer operating systems, like their microcomputer operating system counterparts, are becoming operating environments (see Special Report on Microcomputer Operating Systems, *Computer Design,* July, 1984, p 151). An operating environment contains a rich set of software services provided by the operating system for the minicomputer system designer or software developer. These services are provided while the operating system satisfies a basic set of objectives for the hardware designer. These design objectives include reliability, protection, predictability, convenience, efficiency, general services, flexibility, extensibility, and transparency. Of course, it may be necesssary to compromise some objectives in order to optimize others.

Often, an operating environment allows both minicomputer hardware and application software to call on an operating system for services previously provided by the application program, operating system extensions, specially written software, or dedicated hardware. Lacking any of these, it might not

Operating system design

Minicomputer operating systems can be divided into message passing and procedure-oriented designs. It is only a rough division, because a sound basis for categorization is lacking, and there is a great overlap between the two. Message passing has a small, static number of large processes that share data through message channels. On the other hand, procedure-oriented designs feature a large number of rapidly changing, small processes that communicate through direct data sharing.

As expected, a message-based system must be able to pass messages between processes within the operating system. In contrast, a procedure-oriented system uses system calls to perform its communication chores. Processes cooperate through semaphores, or other software mechanisms such as locks or events.

There are other differences between the two designs. For example, in message passing, process synchronization and resource queuing are implemented in message queues attached to processes. These control the resources in question. In procedure-oriented operating systems, process synchronization and resource queuing are handled by process queues waiting for resource semaphores.

To do its job, the message-passing system assigns process priorities statically when the hardware is designed. The procedure-oriented system provides processes with dynamic priorities. Furthermore, peripheral devices are handled as processes instigated by messages sent to actual devices in message passing. In contrast, procedure-orientation calls for peripherals to be resources used by a process during an I/O operation.

These differences, in the ideal case, should minimally affect the computer system designer's operating system choice. In other words, operating system internals should be transparent to the computer system designer or integrator. With either design system it should be possible to provide all the needed operating system functionality. Although not yet fully developed, the duality theory of operating system design implies that a message-passing system and a procedure-oriented system can accomplish the same functions.

have been possible to provide the service (like program development or fault tolerance) in an efficient way (if at all). For example, in the case of real time or fault tolerance, these two services allow a minicomputer and its hardware to handle such applications as robotics, processs control, and transaction processing at faster, but still reliable, throughput rates.

While the proprietary minicomputer operating systems remain, the trend is clear. And, more will be introduced by such vendors as Digital Equipment Corp (Maynard, Mass), Data General (Westboro, Mass), Perkin-Elmer (Oceanport, NJ), Hewlett-Packard (Cupertino, Calif), and Gould (Santa Clara, Calif). Proprietary minicomputer operating system development is relatively static. Extensions are constantly appearing, but these are more or less minor additions. A really new operating system for minicomputers has not appeared in years. (Note that Unix is about 15 years old.)

These days, end users and computer system designers demand minicomputers that can communicate with their counterparts, and other machines both large and small, in either a standalone or networked operating mode. In short, the days of single-vendor supplied, proprietary operating systems are getting shorter. For many, the need for application software portability, disparate machine-to-machine communication in the future office and factory, and an operating environment that provides a wide variety of software services means that the charted minicomputer operating system course will take one of two tacks.

Either there will be a direct Unix tack under license from AT&T Technologies of Murray Hill, NJ (be it Version 7, System III, System V, or one of the to be announced versions), or there will be Unix compatibility with a licensed Unix derivative or a look-alike (independently developed by various software vendors). Systems in the latter two categories include BSD 4.1 or 4.2 from the University of California at Berkeley, the just-announced NCI Coherent operating system from Network Consulting, Inc (Burnaby, British Columbia, Canada), Pick Systems' (Irvine, Calif) recently announced Release 84 of the Pick operating system, and a range of others.

The big push

Unix will be dominant not because it is an outstanding choice (although it is first-class), but because it is the only one in the race. Unless some major firms introduce major contenders in the next year, the race for the operating system "cup" will be won by default.

As mentioned, existing minicomputer operating systems have a proprietary air to them—and indeed they are geared to a specific manufacturer's hardware. But, AT&T has been very clever in getting Unix to fill an obvious operating system vacuum. In a marketing campaign that puts the lie to those who said that newly liberated AT&T would not know how to compete in a free marketplace, that firm has literally flooded the technical seas with Unix promotion. For example, who has not seen the two-page color ads in every magazine or newspaper geared to the computer designer?

Given the lack of obvious contenders, it would be surprising if Unix did not continue its winning ways in the minicomputer marketplace. Remember, the fine features it offers as an operating system

environment, along with its heavy promotion in academia and industry, have made it way ahead of its time in that regard. Indeed, it was first designed as an environment.

Thus, Unix is already the acknowledged minicomputer leader. Only the extent of the lead and the projected growth figures are debated. Marketing research organizations such as Yates Ventures (Los Altos, Calif), and Gnostic Concepts (Menlo Park, Calif) predict linear Unix growth for several years. The slope of the growth curve is in question though, with Gnostic being the most conservative at 100,000 systems shipped this year. According to Jack Scanlon, vice president of AT&T's Computer Systems Division (Lyle, Ill), the 1982 installed base of Unix operating systems or its derivatives was about 45,000. Today it is about 70,000, and Scanlon expects the figure to double by the end of the year.

New operating system software technology developments promise to make Unix's position even stronger. For example, realtime Unix versions are starting to appear from a variety of vendors. The most notable of these is AT&T, which has bundled such software with its top of the line 3B20D minicomputer. Designed for fault-tolerant applications—the latest computer system design feature—this machine features a symbiotic relationship between duplicated hardware and three-part software (kernel, process, and kernel process). Thus, cost-effective fault-tolerant applications are ensured.

Already used in a variety of telephony-related, fault-tolerant transaction processing applications, the 3B20D represents competition for the previously listed minicomputer manufacturers. These firms, content to let the fault-tolerant market stay with relative old-timer Tandem Computers (Cupertino, Calif), have watched some dozen firms [Auragen Systems Corp (Fort Lee, NJ), August Systems (Tigard, Ore), Stratus Computer (Natick, Mass), and Synapse Computer (Milpitas, Calif), among others] enter the fault-tolerant market with minicomputer hardware and software (often Unix) in the last year.

In fact, the larger minicomputer manufacturers have been strangely idle while the newcomers try to carve out a fault-tolerant market niche. They are so quiet, in fact, that one of the topics discussed at the 14th International Conference on Fault-Tolerant Computing was, "What happens to the new minicomputer vendors if and when the likes of DEC, Data General, and Hewlett-Packard decide to get into the fault-tolerant minicomputer market?"

The computer system designer at the older firms must, of course, come up with an operating system that handles fault-tolerant applications as well as existing customer bases. Of course, designers of new fault-tolerant minicomputer operating systems have no such problem. For their part, the new firms are racing to design hardware and software irresistible enough to computer system integrators, to allow new firms to prosper over the long term.

Until the big boats start sailing, the smaller entries are going about their business combining Unix and fault tolerance. One of them, Auragen Systems, has designed a distributed architecture minicomputer geared to Unix Version 7 fault-tolerant transaction processing applications. To do its job, the firm had to modify the Version 7 kernel. This was to handle the message-passing software it needs to keep backup hardware up-to-date with ongoing software processes. Therefore, backup hardware can perform in the event of a primary hardware failure (see Panel, "Operating system design").

Problem solving

It has been said that everyone likes standards—that is why there are so many of them. Unfortunately, Unix comes in a choice of flavors. As mentioned, there is AT&T licensed System III, Version 7, and System V. All can be made to communicate with each other, but with varying degrees of efficiency. Certainly, before Unix can really "take over," there must be some way to provide this communication until the minicomputer world standardizes.

New operating system software technology developments promise to make Unix's position even stronger.

One way to handle matters until a single Unix standard is formulated is to opt for a minicomputer whose Unix-based operating system can be made to look like AT&T Unix System V or University of California at Berkeley 4.1 BSD. These operating systems are two of the most modern and popular Unixes for minicomputer software developers. Pyramid Technology Corp (Mountain View, Calif) followed this course with its Unix-like operating system resident in the System 90x minicomputer (see *Computer Design* June 15, 1984, p 25). Pyramid, unlike AT&T and Auragen, did not have to make major kernel modifications. (AT&T and Auragen use Unix kernels, while Pyramid uses a Berkeley kernel.) Pyramid's software modifications were mainly in the user interface, partitioned to provide a Berkeley or System V "universe."

The work at AT&T, Auragen, and Pyramid represents what is going on in today's world of minicomputer operating system design to ensure Unix dominance, as well as incorporate fault tolerance and realtime capability into the operating system environment. AT&T is concentrating on real time and fault tolerance with its RTR (Unix) operating system. Auragen (with its Auros operating system) and several other youngsters in the fault-tolerant game are doing likewise. And, as mentioned,

Services for the computer system designer

Harried computer system designers, rushing to satisfy the demand for Unix and realtime operating system capability on their minicomputer systems, need all the help they can get. So, it is fortunate that the Unix craze has spawned another growth period for value added resellers (VARs)—those third parties that take standard manufacturer's minicomputer hardware and add their own software to it. This is done to meet needs that the big hardware folks cannot, or will not meet, because of market size, inability to move quickly, or the need to protect their own software. VAR sales are said to be increasing at rates ranging from 35 to 50 percent per year. Clearly, designers are looking into VARs when they need software that they do not have the time or the expertise to design themselves.

Typical of the many firms supplying value added software services is Uniq Digital Technologies (Batavia, Ill). A DEC distributor (not all VARs are distributors), it ports Unix to DEC's VAX machines. Realizing that the market is undecided as to which Unix variation is best, Uniq offers both AT&T and Berkeley designs. It also supplies driver designs and Unix software tools—the Unify relational data base, for example.

Designers need to examine a VAR's relationship with its large-firm hardware supplier carefully before buying. Some VARs make the large hardware firms nervous—sometimes they modify hardware, make promises to customers that cannot be kept, or do not have adequate resources to deliver—and the hardware firm suffers by extension.

While they cannot be classified as VARs, firms have been spawned to supply computer designers with boards, boxes, and software to integrate Unix and realtime capabilities into minicomputer systems. For example, start-up American Information Systems (Palo Alto, Calif) will soon offer a National Semiconductor 32032-based Xenix-running board for the DEC Q-bus. And among others, Cambridge Digital (Cambridge, Mass) says it will supply Unix System III or V for VAXes and PDPs, and Berkeley 4.2 for VAXes. It also offers realtime Unix for PDPs as does S&H Computer. For its part, Interactive Systems Corp (Santa Monica, Calif)—the single-user Unix supplier for the ubiquitous IBM XT—furnishes its IS/3 Unix for a variety of DEC machines and others.

The services offered by firms cashing in on the operating system software boom cover almost everthing the computer system designer might need to convert to Unix. For example, Rapitech Systems, Inc (New York, NY) will help designers opting for Unix by converting their Fortran programs to Unix's C-code.

Pyramid is into Unix-flavored compatibility (as well as high end minicomputer power).

But, other firms making operating system environment contributions are not Unix related. These tap the market for enhancements to existing, proprietary operating systems. Since it is the leading minicomputer supplier, many firms supply enhanced or upgraded versions of DEC operating systems.

A good example of such a firm is S&H Computer Systems (Nashville, Tenn). For some years now, it has been making modifications to DEC operating systems to turn into operating environments for DEC. It is important for computer system designers and integrators to know about firms like S&H. They supply operating system services that DEC may not find economical to supply. In fact, a variety of firms perform such software chores (see Panel, "Services for the computer system designer"). S&H has progressively updated DEC operating systems for minicomputer applications. Its just released operating system offering allows PDP-11s to run realtime applications.

Getting organized

Operating systems are often referred to as an *ad hoc* collection of software routines to serve whatever purpose is important at the time of their design. There is much truth to this view. Certain industry observers view the current minicomputer concern with real time and fault tolerance as the latest example of operating systems needing to adapt. What, these observers ask, does it take to organize operating system design so it can take care of needed features in an organized way?

Traditionally, as MIT's operating system gurus Stuart E. Madnick and John J. Donovan have put it, the term operating system denotes those program modules within a computer system that govern the control of such equipment resources as processors, main storage, secondary storage, I/O devices, and files. These modules resolve conflicts, optimize performance, simplify the effective use of the computer system, and act as an interface between the user's programs and the physical computer hardware.

This definition, while useful conceptually, does not explain how an operating system can best be designed and organized. In fact, there are few worthwhile answers to the how-to-design question, and the *ad hoc* method seems to dominate the operating system design industry. However, some guidelines aid the computer designer to appreciate the difficulties. Their understanding helps the hardware designer interact with the operating system software designer to ensure that future operating systems meet minicomputer needs.

As might be expected, minicomputer system hardware architecture strongly affects operating system design. For example, most minicomputers are designed around a central memory [Fig 1(a)]. There are exceptions, such as the DEC PDP-11 minicomputer family, which is DEC Unibus oriented [Fig 1(b)]. As might be expected, minicomputer systems in a family vary in speed, memory, and all other features.

Thus, the operating system must accommodate these variations while allowing software portability from model to model. It must take into account

DEADLINE?

Don't get nervous... Get Sudden Service



DIP Sockets IC, ICS, FAB, ICF, ICO, ICC and ICI. Wide variety of contacts & styles.

NEW

Pin Grid Array Sockets Low insertion force of only 3 oz. per pin. Standards with lead counts from 64 to 132 pins. Customs in 8x8 thru 17x17 grid patterns.

NEW

Low Cost .025" Sq. Post Terminal Strips "Snap" to any desired number of positions: 1-36 single row, 2-72 double. Straight, right angle or double row.

Adaptors Stamped-pin with mating covers and machined pin styles.

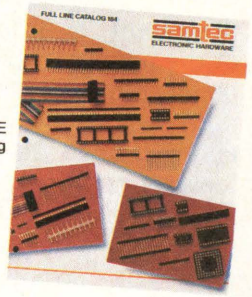
Socket/Terminal Strips Widest variety in the industry. Most with "Snap Strip" breakaway.

NEW

Unique Snap & Lock IDC Strips Locks contact-to-cap at every position...without bulky side locks. 2 thru 72 positions. Mates with .025" square posts.

We've expanded our line... and we're equipped to handle really big orders. But, large or small, you'll still get Samtec Sudden Service... which means you get your order when you need it. I personally guarantee it.

Sam Shine
Prop.



Send for NEW & FREE 52-page catalog

© 1982 Samtec, Inc.

samtec

810 PROGRESS BOULEVARD, NEW ALBANY, INDIANA 47150 (812) 944-6733

CIRCLE 83

differences in word size, asynchronous and synchronous operation, such different register types as general purpose, stack pointer and program counter, multiprogramming, vectored or nonvectored interrupts, a variety of instruction sets, memory management and floating point arithmetic capabilities, and more. Accommodating these features is difficult enough in microcomputer operating systems and environments; it is far harder on minicomputers.

Once the minicomputer operating system can handle these hardware variations within a minicomputer family, the question arises as to what services it should provide. Here, the operating system and computer system designers or integrators have different points of view that can be summarized (not totally tongue-in-cheek) as the operating system designer saying that a request is impossible to satisfy, and the computer system designer saying that the system will not meet customer needs unless it includes that particular feature. As with all engineering problems, the answer is arrived at by compromise and mutually acceptable criteria. Criteria might include discussion of such objectives as the degree to which batch processing is accommodated. This old style of computer operation is still furnished by the latest minicomputer operating systems. For example, both Data General's AOS and DEC's VMS operating systems handle batch processing.

Defining terms and needs

Hardware and software designers must also agree on the degree to which realtime capability is furnished. Real time is a buzzword that means different things to different people. Real time to a telephone user querying a minicomputer that provides 800 numbers is not the same as real time to pilots depending on an embedded computer-based weapons control system to protect their F-15 from enemy fire. Though most minicomputer vendors provide a realtime executive—(a scaled-down operating system with minimal services), even these are turning more environment-like as service demands grow.

Then, there is the need for multiprocessing, multiprogramming, time-sharing, and networking. The minicomputer system designer requires an operating system able to handle a large number of microprocessors. These are thrown at all computer system peripherals in order to make them intelligent—all in the name of off-loading the processor so it can be most efficient at its computing chores. Operating systems must also handle the array processors and pipelined devices that are added to today's loosely or tightly coupled minicomputers. DEC's just introduced VAX-11/785, running both the VMS and the Ultrix-32 operating system, is the latest multiprocessing minicomputer. Data General's recent top of the line Eclipse MV-1000 is also a multiprocessing machine with its own special operating systems.

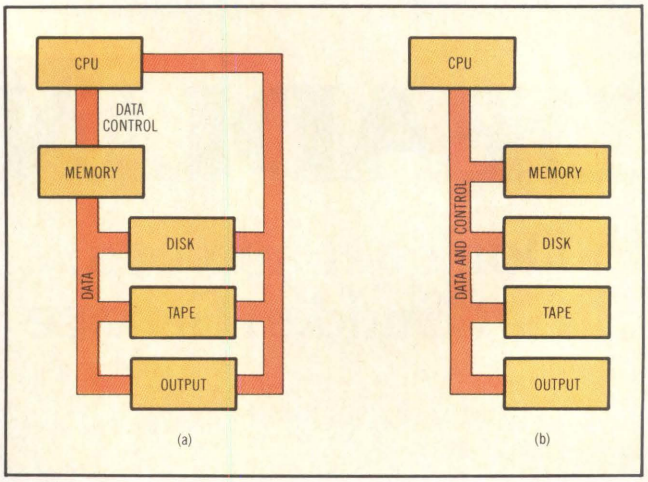


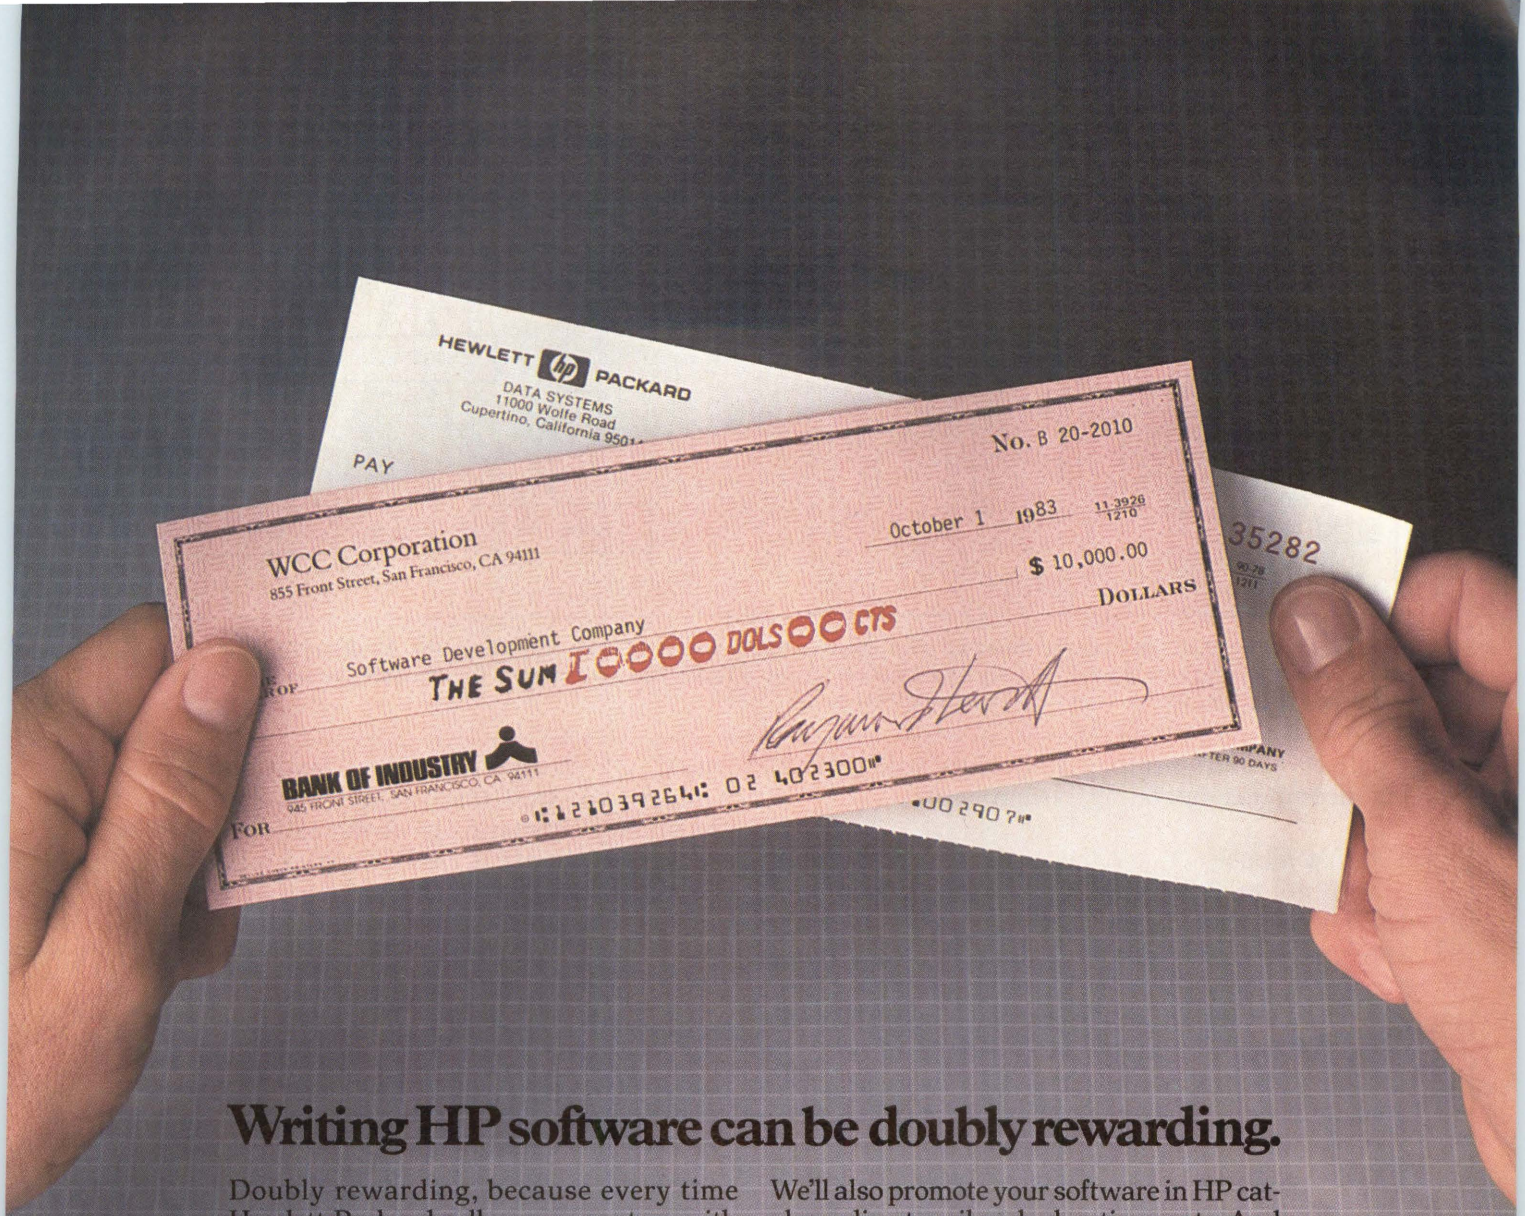
Fig 1 Most minicomputer architectures sport a central memory to connect a CPU and its peripherals (a). With some exceptions, the minicomputer operating system must be designed around this approach. DEC's PDP-11 features a central bus for data and control (b).

Multiprogramming, of course, is related to the concurrency that is so much in demand to execute more than one program at a time. Concurrency is even in demand in microcomputers (see *Computer Design,* July 1984, p 187). Time-sharing and concurrency are related. Time-sharing lets many use all but one processor, and also requires an operating system that takes care of the timing, priority setting, and scheduling that multiprogramming demands.

Finally, there is minicomputer networking. The fact that machines (regardless of vendor) are hooked together these days puts special demands on the operating system. For example, if the network allows disparate machines from disparate manufacturers to communicate fully, it needs to handle, either through the operating system or a frontend processor, the International Standards Organization's (ISO) seven-layer protocol for computer communications (*Computer Design,* June 15, 1984, p 57).

It also must be agreed that an operating system should have such minicomputer attributes as processor time, memory management, peripheral devices, software resources, realtime applications, and fault tolerance. But, deciding on the right balance is a major problem.

For example, since many minicomputer calculations are I/O bound, how does an operating system allocate its processor time as a resource? Since for cost's sake, the minicomputer system must be shared among multi-users, what is the best way to do this? Is the best design run-to-completion, a fixed or flexible priority scheme? Memory management presents similar problems. Just as processor time must be scheduled, so must memory, since programs cannot execute outside of main memory. While operating systems boast a wide variety of memory management schemes, and ever-cheaper RAM is making



Writing HP software can be doubly rewarding.

Doubly rewarding, because every time Hewlett-Packard sells a new system with your software, you'll not only get a check from the customer, but one from HP as well!

As thanks for writing software for our systems, we'll give you a bonus of 30% of your software's sales price—up to 6% of the net HP system's price. That could mean $3,000 on a $50,000 system... or as much as $30,000 on a $500,000 system. And you'll get it *every time a new system is sold with your software.*

We won't tell you what kind of program to write. What industry to write it for. Or even specify the system. Our only restriction is that, to qualify for the bonus, your software must sell for at least $10,000.

Of course, extra cash isn't the only incentive. To help you get started, we'll sell you a development computer at a 40% discount.

We'll also promote your software in HP catalogs, direct mail and advertisements. And you'll have the benefits of being associated with HP—one of the worldwide leaders in computing systems. A company whose products range from widely-used business computers, like the HP 3000, to one of the world's most advanced 32-bit computers, the HP 9000.

If you'd like to learn more, write to Hewlett-Packard, Attn Gwen Miller, Dept. 11173, 19447 Pruneridge Avenue, Cupertino, CA 95014. In Europe, write to Henk van Lammeren, Hewlett-Packard, Dept. 11173, P.O. Box 529, 1180 AM Amstelveen, The Netherlands.

We'll give you all the details on our HP Plus Software Supplier program. And all the reasons why it will be rewarding for *you* to develop software for HP.

HEWLETT PACKARD

U.S. List prices only.
BD-02301

CIRCLE 84

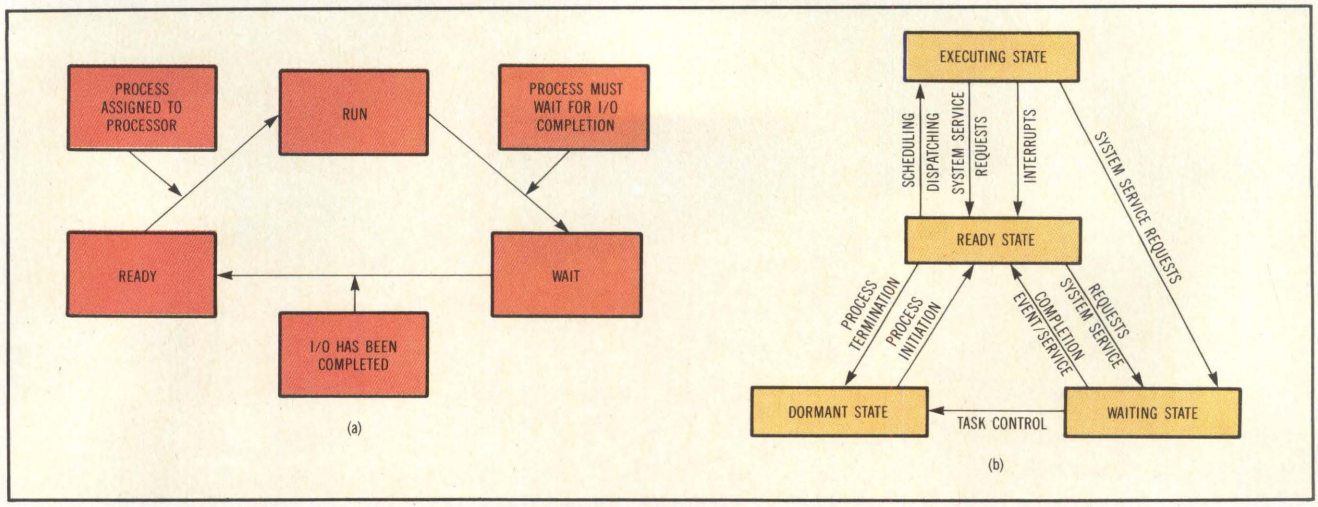


Fig 2 Processes can have simple states (a) or a more complicated set of state transitions (b). Realtime minicomputer designers should look for the simplest set for a really fast system and, for example, avoid wait states with memory, file, device, and other wait subdivisions.

matters easier, memory control is still a major minicomputer operating system chore.

Getting formal

It is not easy to formally design a minicomputer. In this design approach, a formal specification is set up *a priori*, rather than as a collection of software features to satisfy whatever is "hot" at the moment. Tied in with the formal operating system specification concept is the abstract machine notion. An abstract machine maps whatever software functions the minicomputer user wants into specific hardware operations. In fact, in an extension of the concept developed by famed computer scientist Edward W. Dijkstra (he invented the semaphore concept and its use in synchronizing cooperating software processes), each operating system layer represents an abstract machine built on the lower layer machines, and extending the operating environment for the minicomputer user's convenience. Key to this abstract model is the concept of an " object" or implemented software feature(s).

The object approach lets the computer environment be easily modified to new applications.

An object comprises both the memory that contains the object's descriptors and the software methods that are used to manipulate it. Methods related to procedures in conventional operating systems specify the operations that are performed on an object. To communicate, objects send messages to each other.

Alas, minicomputer operating system design is still far from the ideal case of formal operating systems based on abstract machines and objects—and remains *ad hoc*. One notable exception is the Smalltalk-80 operating environment out of Xerox's Research Center (Palo Alto, Calif). Further research is needed to develop an understanding of abstract machine properties for formally specifying functions. Xerox has made a start, though, even if Smalltalk-80 is used mostly by the computer science community.

The greatest advantage of the object approach is that it allows the computer environment to be easily modified to new applications. Most observers expect the new, flexible 32-bit microprocessor-based workstations for software applicaton development to stimulate object-oriented operating system use. These machines furnish such operating systems with the powerful hardware and graphics needed to be cost effective.

Xerox is not alone in its object-orientation. For example, Telenova, Inc (Los Gatos, Calif) has used similar concepts in its voice/data handling operating system design (*Computer Design*, July 1984, p 231) as have a few other firms.

Another approach, closely related to the abstract machine but with some implementation success—mostly at IBM (Armonk, NY)—is the virtual machine. A virtual machine is implemented through a virtual machine monitor (VMM). The VMM is a software program that executes on real hardware and provides a set of virtual resources (which may or may not be physically present). The purpose of a virtual machine is to provide the user or users of minicomputer (or mainframe) with one or more environments. Thus, under VMM control, different users call different operating systems and environments. While some work has been accomplished in this area, there is not too much demand. And worse, most real computers cannot support the concept.

Most practical minicomputer operating systems (such as those discussed) depend upon either the programmed operator design or the monitor design. In the former, all operating system functions are extensions of the user's programming language provided as modules (which do not interact). The operating system merely responds to user requests. Most low end minicomputer operating systems are designed this way, as are microcomputer operating systems like CP/M and PC-DOS.

Various techniques

In the contrasting monitor design approach, the minicomputer operating system has a collection of procedures and data structures that can be shared among software processes, albeit one process at a time. The monitor approach is similar in concept to the object-oriented design. If a requesting process finds the needed procedure or data structure in use, it may have to wait in a queue.

In other words, the operating system is a resource collection protected by the monitor. With multiple resource copies, each has a reentrant monitor that allows each one of them to be active. Dijkstra's previously mentioned semaphores usually implement monitors.

Of course, a variety of operating system design techniques can be brought to bear that do not fall into neatly generalized categories. Many of them are clever design concepts. For example, in the design of realtime operating systems, it is valuable to simplify process states in order to achieve the fastest possible scheduling [Fig 2(a)].

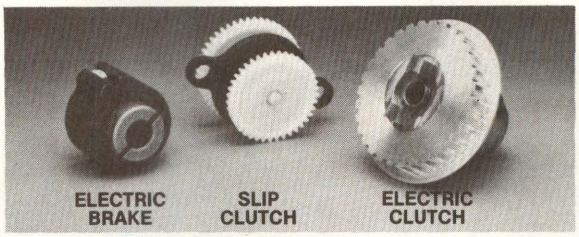
In fact, processes may exist in many states in a typical minicomputer operating system [Fig 2(b)]. Each process makes state transitions in response to instructions, other processes, or peripheral devices. Clearly, faster transitions follow from the simplest state structure. This structure does not allow multiple wait states that must be identified and communicated with (by means of software overhead).

The latest operating system design thinking calls for the monitor concept to be enhanced by the so-called manager. A manager, thought of as an enhanced object-oriented software entity, combines the ideas of a monitor and a process, allowing service requests to be satisfied immediately. Blocking occurs only when inputs or responses must be received.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 722 *Average 723* *Low 724*

RPM custom-designed clutches and brakes... The cost-effective alternative for your rotary motion problems



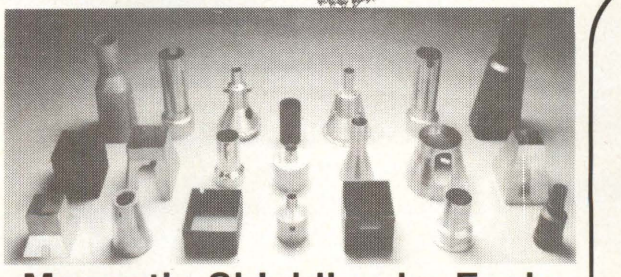
RPM clutches and brakes have solved many rotary motion problems for OEMs. That is our specialty. We design and manufacture our products to meet your special performance needs. Of course, we also offer standard units.

EXCELLENCE rpm
REELL PRECISION MANUFACTURING CORPORATION
1259 Wolters Boulevard
St. Paul, Minnesota 55110

CALL FOR FREE APPLICATION ANALYSIS (612) 484-2447
RPM offers free analysis of your OEM clutch and brake applications requiring up to 50 in-lbs. of torque.

CIRCLE 85

Excellence In It's Field



Magnetic Shielding by Eagle

CUSTOM SHIELDS • STANDARD SHIELDS • FINISHING • DESIGN ENGINEERING • FOIL • HEAT TREATING • SHEET • TESTING • CONSULTING • FABRICATION

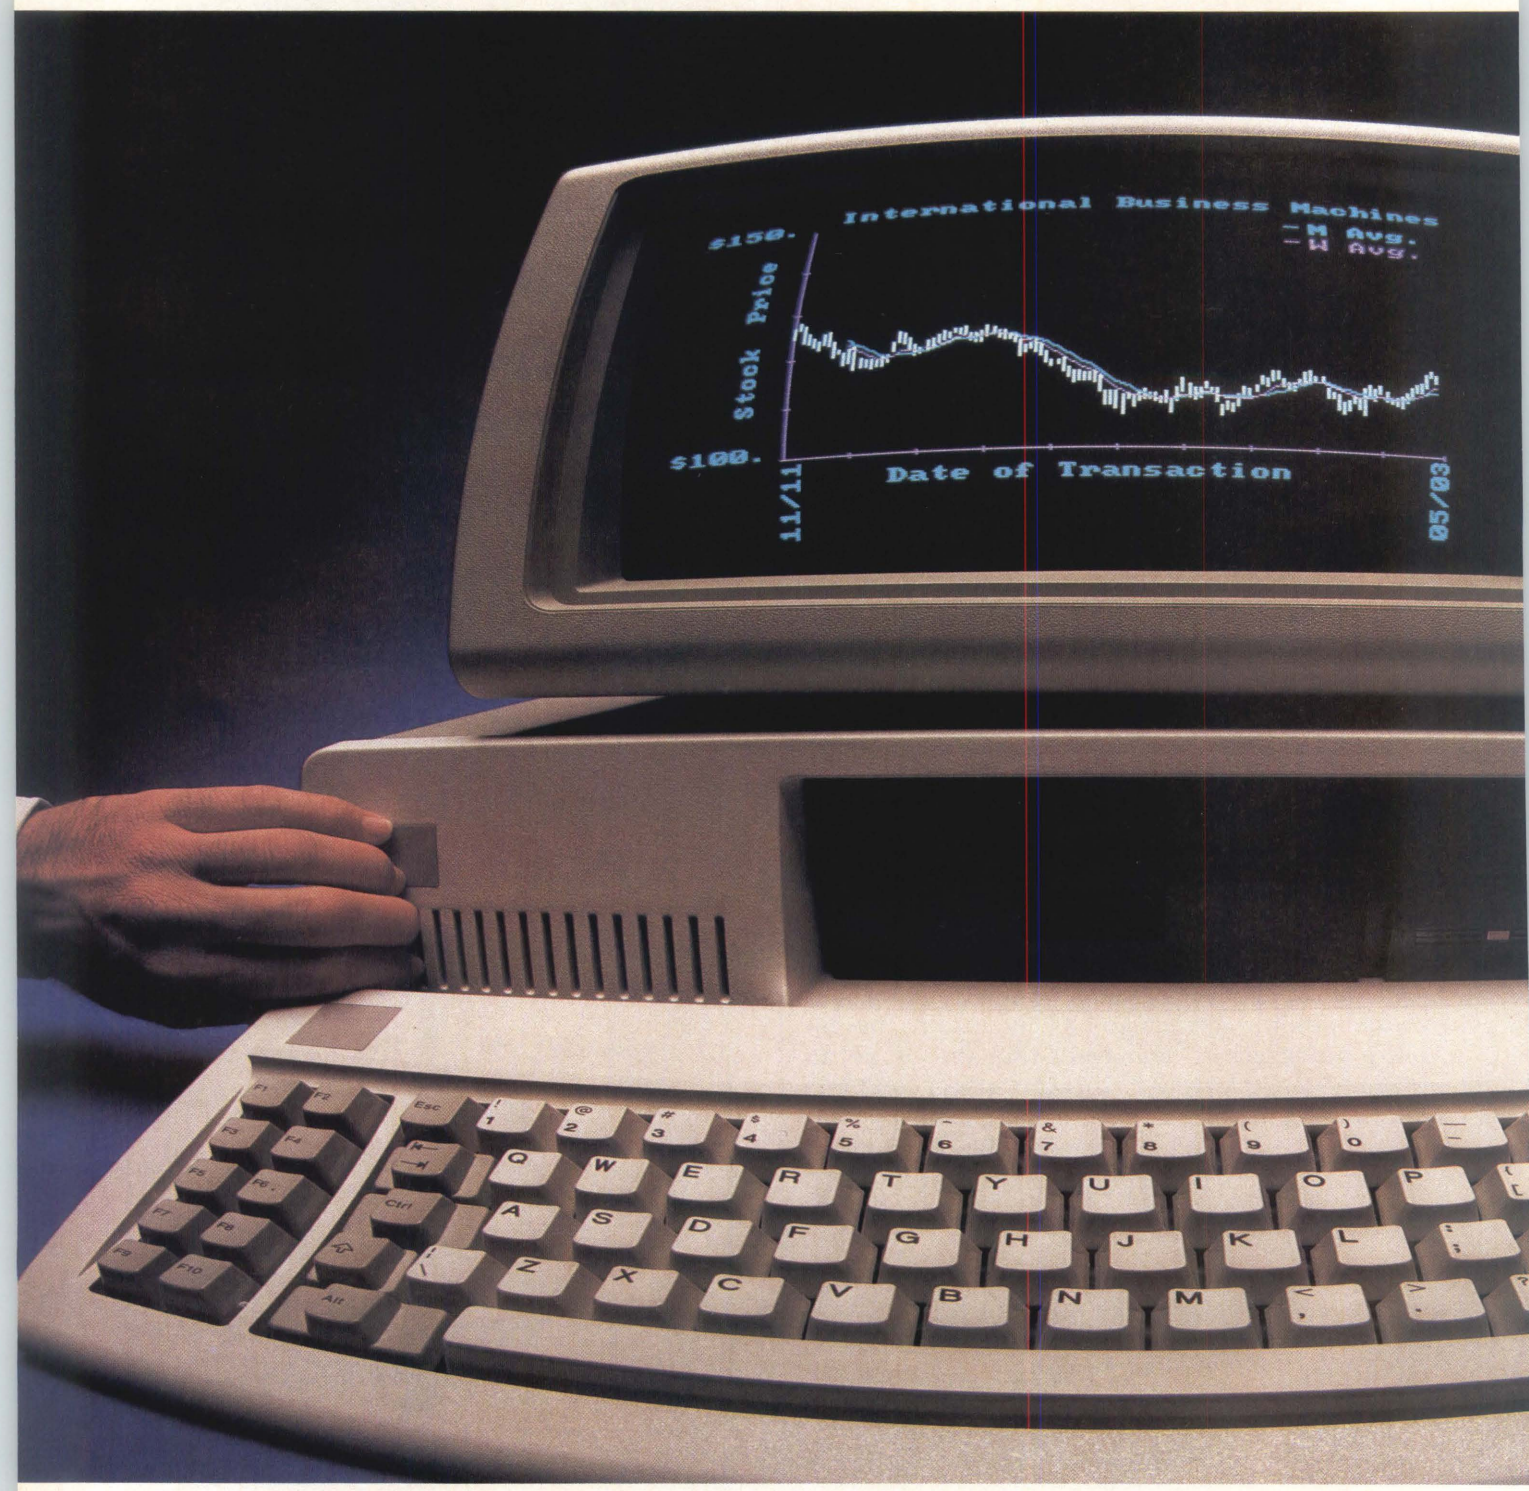
Eagle can help improve your product, and lower costs, by designing the right shield for you. Take advantage of Eagle's vast background in shield design and production. Choose from a wide selection of sheet and foil, so you can form your own shields. For helpful design and cost data, write or call. Offices worldwide.

EAGLE MAGNETIC COMPANY, INC.

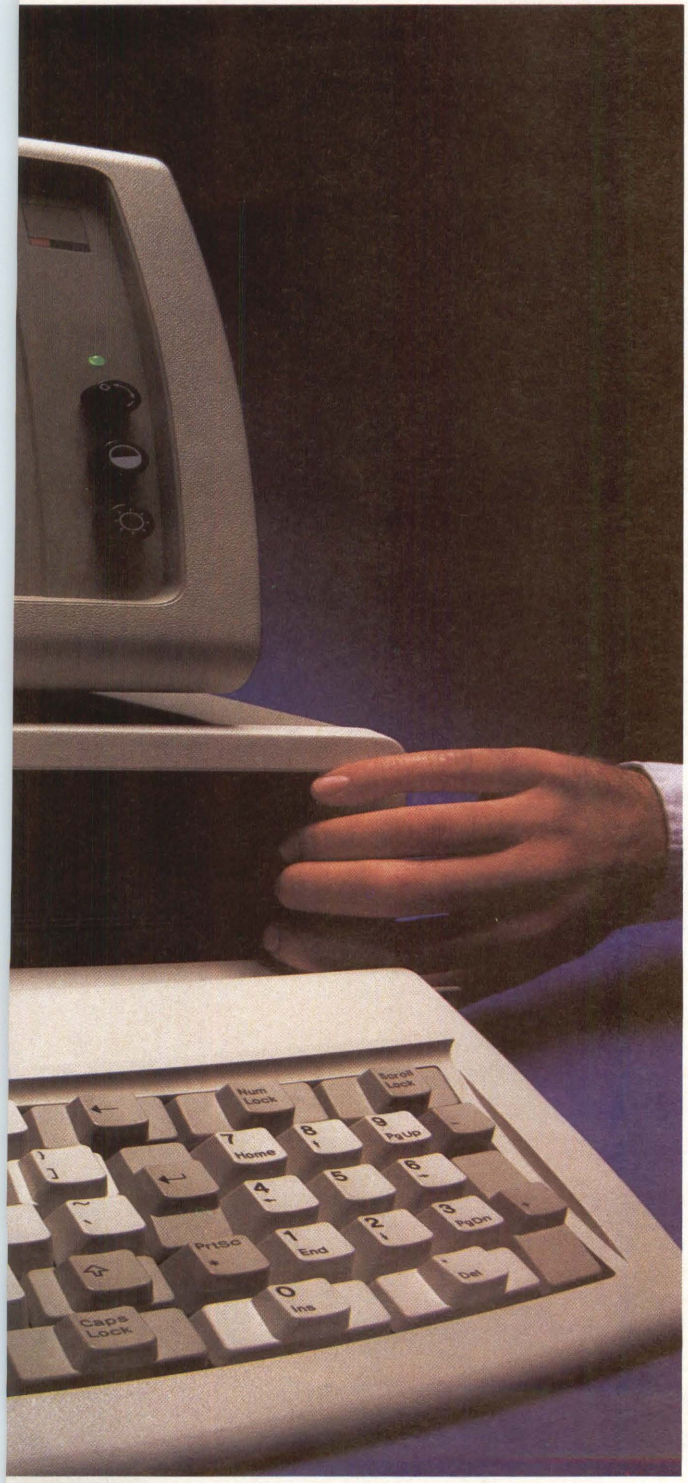
P.O. BOX 24283 • INDIANAPOLIS, INDIANA, 46224 • PHONE (317)297-1030

CIRCLE 86

HOW TO STRETCH A PC TO CREATE NEW MARKETS.



OR EXPAND EXISTING ONES.



When is a PC not a PC? When you turn it into a personal data base cruncher. Or a personal graphics workstation. A personal image processor. Even a personal signal processor.

You can do all that and more with the Marinco board pictured below. It's your key to unlocking new opportunities and new markets.

Power to the PC.

We've put array processor technology to work to supercharge the PC. Suddenly a spectrum of applications that were beyond that PC's capabilities are wide-open possibilities.

If your strength is software, the Marinco array processor gives you incredible performance improvements. Even the most complex spreadsheets can ripple faster than you can say 1-2-3. Enormous data bases can search, sort and report in the blink of a cursor. And high speed arithmetic calculations which slow a standard PC down to a crawl are performed in split seconds.

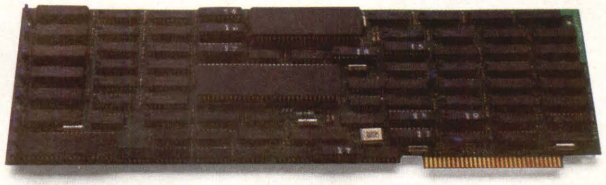
If you're a systems integrator, VAR or OEM looking for a competitive advantage, a PC with a Marinco array processor is a natural line extension, enabling you to offer a cost-effective "personal" version of your best-selling product.

The Race Is to the Swift.

The Marinco array processor is especially flexible and easy to program. So you can get to market quickly. Which, after all, is the name of the game.

We could fill pages with pertinent performance specs on our PC-compatible array processor. So we did. We've put together an 8-page brochure that will stretch your imagination. And just might open a world of opportunity. Write for a copy. Marinco Computer Products, 3878 Ruffin Rd., San Diego, CA 92123. Or call 1-800-421-4807, in CA (619) 268-4814.

Marinco array processors include a 24-bit floating point board with optional IEEE-754 compatibility and a 16-bit integer-only board for maximum price/performance.



MARINCO

CIRCLE 87

When you're designing a first-rate computer system, it doesn't make sense to compromise with a second-rate monitor.

Consider PGS instead. We make three no-compromise monitors, all fully compatible with the IBM-PC, to match your requirements and your budget.

High resolution PGS set the price/performance standard for high resolution RGB color monitors with the HX-12: .31mm dot pitch, 690 dots horizontal resolution, and precise color convergence for a crisp, sharp image.

Super resolution And, when your specifications call for super resolution, there's our no-compromise SR-12: an RGB monitor with a horizontal scan rate of 27.5 KHz which supports 690x480 resolution in non-interlaced mode.

Monochrome For price/performance in a monochrome monitor, we've set the standard with the MAX-12: our amber monitor with dynamic focusing circuitry which ensures sharpness not only in the center but also in the edges and corners.

The monitor to meet your needs

All three PGS monitors are engineered for no-compromise performance to offer you a cleaner, sharper image than any other monitor in the same price class. The HX-12 and the SR-12 both feature uncompromising color convergence for crisp whites without color bleed. The MAX-12 offers impressive clarity in an amber phosphor monitor that runs off a standard monochrome card—no special card is required.

At PGS, our no-compromise approach includes all the details, too, from non-glare screen to a shielded cable—standard features on all PGS monitors, color or monochrome.

Call us at 800-221-1490

Compare your specifications to ours, listed below. Then call us at 800-221-1490 and we'll send you a fully detailed spec sheet plus everything else you need to know about all three no-compromise PGS monitors.

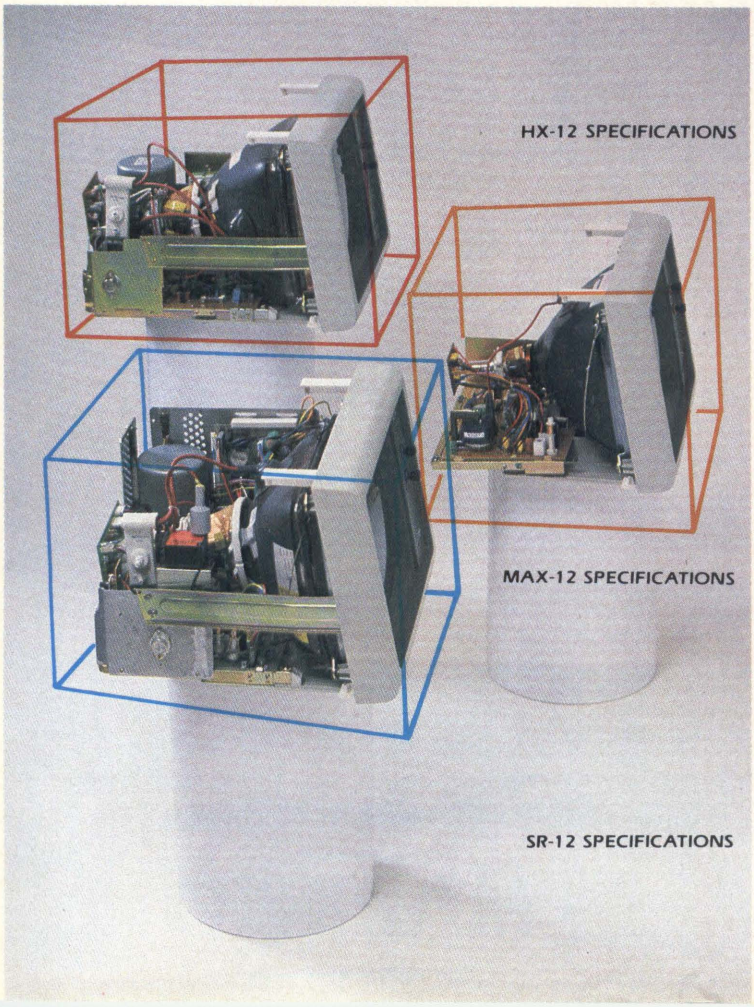
Don't compromise. Look to PGS for the image your eyes deserve.

PGS **Princeton Graphic Systems**

1101-I State Road
Princeton NJ 08540
(609) 683-1660
TLX 6857009 PGS Prin

Nationwide service provided by Bell and Howell Service Company and MAI Sorbus Service Company.

PGS delivers 3 no-compromise ways to improve your image



HX-12 SPECIFICATIONS

CRT	12" Diagonal, 76 degree, In-Line Gun, .31mm dot pitch black matrix, non-glare surface (NEC 320CGB22)
Input Signals	R, G, B, channels, Horz Sync, Vert Sync, Intensity—all positive TTL levels
Video Bandwidth	15 MHz
Scan Frequencies	Horizontal: 15.75 KHz Vertical: 60 Hz
Display Size	215mm x 160mm
Resolution	Horizontal: 690 dots Vertical: 240 lines (non-interlaced) 480 lines (interlaced)
Misconvergence	Center: .6mm max Corner: 1.1mm max
Display Colors	16 colors (black, blue, green, cyan, red, magenta, yellow, white, each with 2 intensity levels)
Characters	2000 characters (80 characters x 25 rows—8x8 dots)
Input Connector	9 Pin (DB9)—cable supplied to plug directly to IBM PC

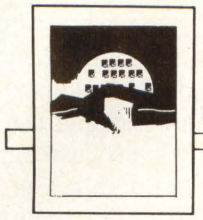
MAX-12 SPECIFICATIONS

CRT	12" Diagonal, 90 Degree, non-glare surface (P 34 Phosphor)
Input Signals	Video signal, Horz Sync, Intensity—positive TTL levels, Vertical Sync—negative TTL levels
Video bandwidth	18 MHz
Scan frequencies	Horizontal: 18.432 KHz Vertical: 50 Hz
Display size	204mm x 135mm
Resolution	Horizontal: 900 dots Vertical: 350 lines
Input Connector	9 Pin (DB9)—cable supplied to plug directly to IBM PC

SR-12 SPECIFICATIONS

CRT	12" Diagonal, 90 Degree, In-Line Gun, .31mm dot pitch black matrix, non-glare surface
Input Signals	R, G, B channels, Horz Sync, Vert Sync, Intensity—all positive TTL levels
Video bandwidth	25 MHz
Scan frequencies	Horizontal: 31.5 KHz Vertical: 60 Hz
Display size	215mm x 160mm
Resolution	Horizontal: 690 dots Vertical: 480 lines (non-interlaced)
Misconvergence	Center: .5mm max Corner: 1.0mm max
Display colors	16 colors (black, blue, green, cyan, red, magenta, yellow, white, each with 2 intensity levels)
Characters	2000 characters (80 characters x 25 rows)
Input Connector	9 Pin (DB9)—cable supplied

CIRCLE 88



OPERATING SYSTEM FEATURES REAL TIME AND FAULT TOLERANCE

A third Unix operating system layer, known as the kernel process, is the key software design innovation.

by Bob Snead, Frank Ho, and Bob Engram

The combination of a realtime operating system and fault-tolerant environment furnishes new capabilities for the minicomputer operating system integrator. A system that operates continuously during diagnostics, repair, maintenance, software updates, and system administration would be ideal in such computer system design and integration applications as banking, process control, and military operations. One such system, the AT&T Bell Labs Unix RTR and 3B20D minicomputer, is already used extensively in telecommunications switching networks. This operating system provides a realtime, high availability Unix operating system environment (see Panel, "A wide range of applications").

Unix RTR's architecture is designed to maximize the effectiveness of the underlying reliable hardware (Fig 1), protect the system against software faults, provide support for realtime applications, and provide the Unix process environment. Unix operating system software is usually divided into a process layer and a kernel layer. Unix RTR, however, adds a third layer called the kernel process layer (Fig 2). This extra layer gives Unix RTR its realtime capabilities and establishes the design framework for reliable operating system and application software.

The process layer supports the standard Unix operating system environment. To take advantage of the large physical address space of the 3B20D, each process can comprise up to 128 segments, each from 1 to 128 Kbytes in length. Although processes are swapped, in Unix RTR, certain processes can be made "non-swappable" if the latency caused by swapping is unacceptable.

Code sharing and data sharing are available through the use of a flexible, shared library facility. Defining such a library is much the same as defining processes, except that the library never runs by itself. Instead, any process that needs access to the code or data within simply specifies the inclusion of the library in its address space. All text segments in the library are always shared; data, on the other hand, can be either shared or private. The data in a library is usually shared. However, if the text within the

Bob Snead is a member of the technical staff at AT&T Bell Labs, 1100 E Warrenville Rd, Naperville, IL 60566. He holds a BS and MS in computer science and an MA in psychology from the University of Kansas.

Frank Ho is a supervisor at AT&T Bell Labs. He holds BS, MS and PhD degrees in computer science from the University of California at Berkeley.

Bob Engram is department head of operating systems development at AT&T Bell Labs. He holds a BS in electrical engineering from Howard University and an MS in electrical engineering from Stanford University.

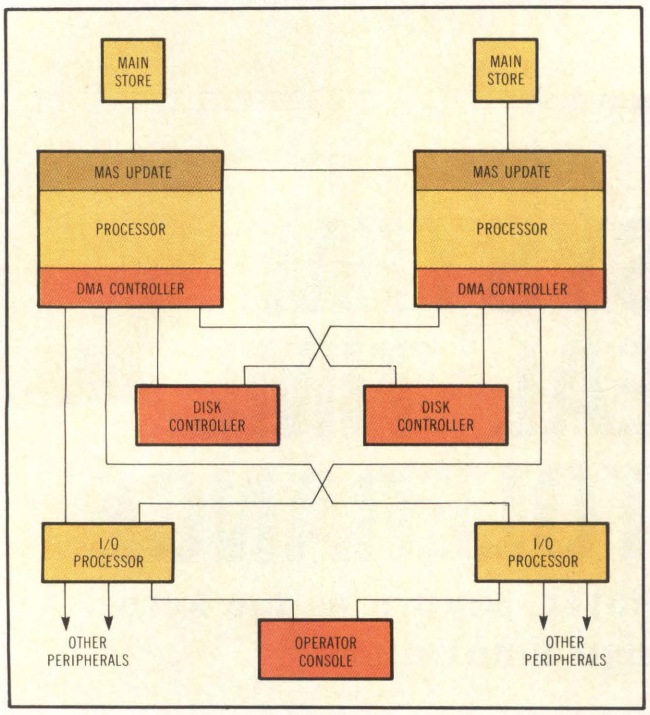


Fig 1 The 3B20D minicomputer hardware architecture is designed for fault tolerance. It features duplicated hardware for critical components. Applications for the computer system designer or integrator include transaction processing, process control, and aerospace.

Unix RTR library requires separate data areas for each process, a fresh copy of those data areas is created. This happens whenever a process that includes the library is created.

The kernel process layer functions to provide a process environment with realtime attributes. One class of processes in this layer features the device drivers that are part of the Unix kernel. Note that in Unix RTR, each driver is given its own address space, thus making it easy to add new drivers, and enhance their reliability by isolating them. Each of these kernel processes can have up to 128 segments and can use shared libraries (even sharing with Unix processes), but each is locked in memory (no swapping).

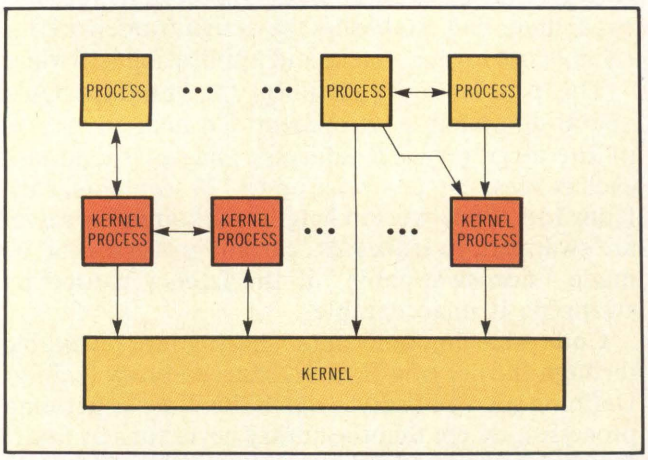


Fig 2 The software architecture of the Unix RTR realtime operating system in the 3B20D minicomputer is comprised of three layers: process, kernel, and kernel process. The kernel process is new and is not used in standard Unix. Processes and kernel processes can communicate with each other as shown.

Control over real time is provided via 16 execution levels which form a static, preemptive priority structure. Unix processes run at the lowest two execution levels (0 and 1), certain system processes (memory manager, scheduler, etc) run at level 2, and the remaining levels are used by kernel processes. Each kernel process runs at a predetermined execution level. By assessing the relative realtime needs of each kernel process, an application can be engineered by placing those kernel processes that require lower latency at higher execution levels.

A set of mechanisms facilitates interprocess communication. For example, there is a general message-passing mechanism that can be used by any process to communicate with any other process. Unix processes can communicate with kernel processes using a more efficient synchronous trapping mechanism.

Managing duplicate hardware and overload

The error interrupt handler (EIH) kernel process monitors faults caused by both hardware and software (see Panel, "A reliable hardware architecture"). When hardware faults occur, the appropriate driver is notified. When software faults occur, the software in question is notified.

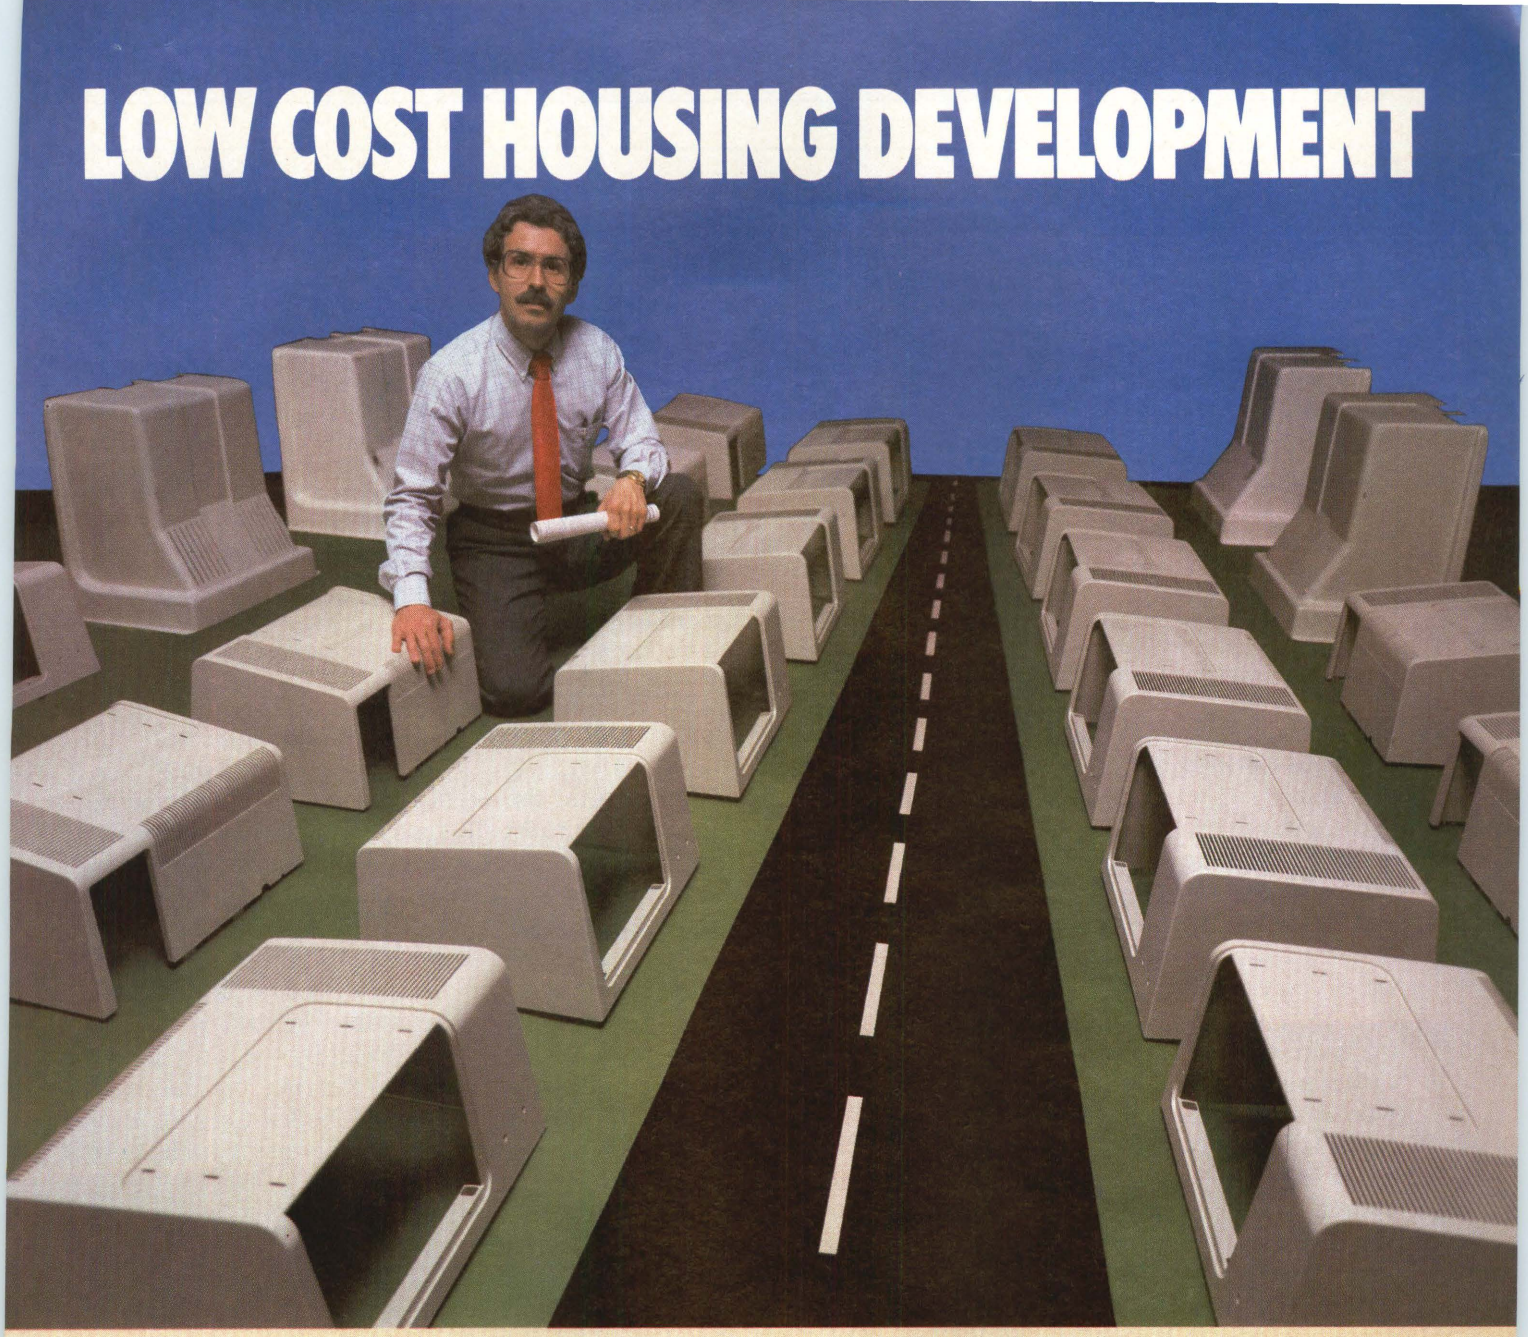
All device drivers share a library that manages hardware device configuration. This library contains routines that keep track of error counts and rates of all the hardware units. When a driver is notified by EIH that a fault has occurred in one of its units, it records this fact. If that unit has become too faulty, then the library changes the configuration by switching to the other unit in a duplicated pair. If EIH detects that numerous faults are occurring at a high rate, or that reconfigurations have not prevented faults, it requests a system initialization.

Exhaustion of one or more system resources during a temporary or prolonged period may, in ordinary systems, cause the system to go down. However, Unix RTR monitors such critical system resources so that overload periods can be detected and corrective action taken without removing the system from service. Among the resources monitored are main memory, CPU, and numerous software resources such as message buffers.

For system message buffers, a progressive strategy is used to recover. First, noncritical "hogs" of the resources are given high priority so that they can release the buffers. If that strategy is ineffective, then hogs are terminated. Finally, if necessary, the system will go through initialization.

Software errors can still occasionally cause data corruption, even with testing. This problem is

LOW COST HOUSING DEVELOPMENT



HOSTYREN™ VO Polystyrenes can help you build important savings and high performance properties into your enclosures.

It's easy to see why more and more designers and processors are switching to American Hoechst's versatile line of HOSTYREN VO polystyrenes for many kinds of electronic enclosures.

Unlike competitive materials, these easy-to-process VO resins give you all the high performance characteristics you need—*without costly over-engineering!*

As a cost-effective alternative to traditional materials, they have a lot to offer: proven color stability, excellent heat and impact resistance, good finish properties and dependable product consistency, resulting from total "in-house" production control.

The complete HOSTYREN polystyrene family is also the *only* line of polystyrenes that offers a choice of custom and standard colors on *less-than-truckload* orders.

Add the fact that American Hoechst's experienced and fast-growing technical sales and service team offers prompt, hands-on solutions to on-site problems and you can see why HOSTYREN polystyrene resins are such a popular choice. Mail the coupon now for technical details and cost comparison data. We're the Plastics Fanatics.

☐ Please send technical information and cost comparison data sheets on your HOSTYREN VO polystyrene resins.

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Telephone _____

CD-8

Send to: Mr. Larry Vallon,
American Hoechst Corporation
289 North Main Street
Leominster, MA 01453 Phone (617) 534-2627

Hoechst

*The name and logo Hoechst are registered trademarks of Hoechst AG.

CIRCLE 89

A wide range of applications

The 3B20D and Unix RTR were both originally designed as the central processor and operating system for several telephony-related applications. All such applications have very stringent availability requirements. The realtime environment provided by Unix RTR is the keystone for several of these applications, such as a telephone operator console system called the Traffic Service Position System (TSPS) and AT&T's newest switch, the 5ESS digital switching system. Several other applications use the machine as a database system. For example, the Network Control Point (NCP) currently handles 800 and INWATS service, while the Attached Processor System (APS) handles the database end of a 4ESS toll digital switch.

Other than telephony, many applications could profit from either the realtime environment or the high availability (or both) offered by the 3B20D/Unix RTR system. A banking transaction management system, for example, would find this system ideal since downtime means money lost. Control applications could benefit from both the realtime environment and the high availability. An air traffic control system would need the realtime responsiveness of the kernel process environment, and depend upon the reliability to avoid the possibly devastating effects of system outage.

By the end of February 1984 there were 231 3B20D/Unix RTR systems being used in telephony applications. Since the first deployments in early 1981, the amount of system outage time has decreased as projected by the system's designers. Thus, in 1981 there were 100 min downtime per year per system, about 50 in 1982, and 25 in 1983. The current per-system average downtime is about 10 min per year per system.

intensified in a system that remains up over greatly extended periods of time. When the system goes through a software initialization, the problem is made even worse since the highly asynchronous nature of software initialization can occasionally leave system structures in an inconsistent state. But, Unix RTR has a number of software audits that detect and often correct these situations. For example, its system-wide segment table can be audited—up to one percent of system time can be spent examining system tables and other data structures.

Occasionally, a process may incur a transient software fault. The system philosophy for handling such faults is to allow the process to correct the problem, or at least to clean it up, rather than imposing a system-wide stategy. This design allows critical system processes to take the recovery action appropriate for its service, and still remain on the air. As might be expected, the overall system integrity is constantly monitored. Such things as overload status, audit results, and hardware indicators are constantly interpreted by a monitor process. If it is judged appropriate, the monitor will automatically start an initialization.

Simply put, the overall philosophy of system initialization is to let the recovery action fit the problem. Thus, the system has five initialization levels. The lowest is a software initialization for applications; next is software for the operating system. More critical still is an initialization with a boot but with preservation of memory resident data bases. The next level involves reloading the data bases with a boot. Finally, this is followed by the most severe level, a boot with a memory clearing and database loading. The system starts at the least severe level and tries it four times before progressing to the next initialization level.

Minimizing downtime

Most computer systems schedule maintenance downtime to update system software or hardware, but downtime is unacceptable for applications that must be available continuously. Unix RTR minimizes system downtime by accepting software

A reliable hardware architecture

High availability is one of the major objectives in the design of the 3B20D processor. The hardware architecture duplicates every major component. Duplication allows the system to remain operational in the event of all single and some multiple faults.

The processor is a 32-bit machine with a 24-bit virtual and physical addressing capability. It incorporates extensive self-checking logic to allow immediate error detection and recovery without invoking diagnostic programs to identify the faulty units. The processors run in active/standby mode, so that a failure in the active processor will cause an immediate switch to the standby processor. The processor memories are updated by the memory update unit (concurrent with instruction execution).

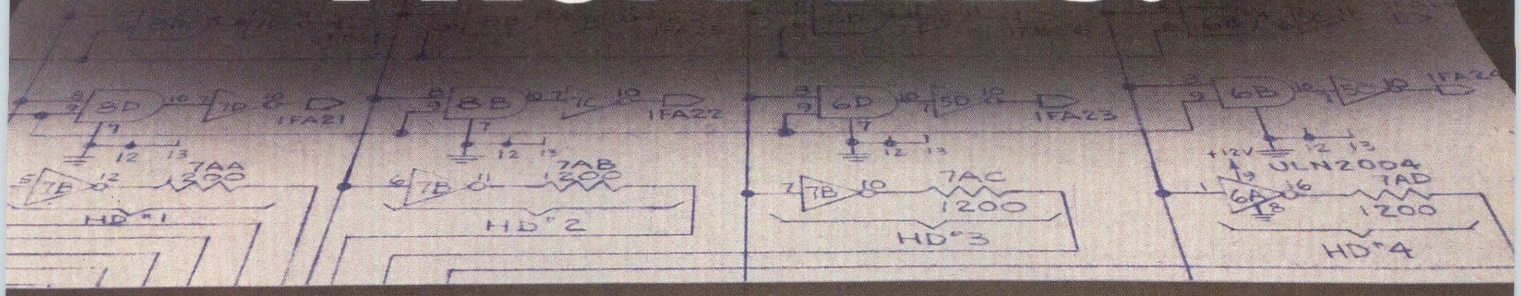
Disk controllers and I/O processors are duplicated and are dual ported to both processors. In addition, disks are duplicated, with disks in a duplicated pair connected to different disk controllers.

Writes to a file cause the disk driver to issue two writes, one to each of the disks in a duplicated disk pair. File reads are translated into disk reads which are alternated between the two drives to reduce access time.

Critical peripheral devices (eg, the operator console), are dual ported and are connected to two I/O processors. Thus, the operator always has direct access to control the system. As a result, failures in a processor, a disk controller or a disk, or an I/O processor will not affect system operation.

GET THIS.

FROM THIS.



MULTIWIRE® DESIGN CENTERS PROVIDE SCHEMATIC REDUCTION.

Working with a sophisticated CAD system, your Multiwire Design Center can shorten your design cycle by weeks, saving you thousands of dollars.

We offer a variety of design services, depending upon your needs. Your boards are designed from as little input as a logic diagram or net list, along with board geometry and special parameters.

Multiwire designs and fabricates finished boards to meet both your routine and accelerated delivery requirements. In fact, you can receive a finished board from schematic in as little as four to six weeks, including design time.

And when you need to make changes, they will be simple and fast. That's because Multiwire boards only use artwork to distribute the power and ground potentials. So a logic change is simply keyed into the computer, creating new floppies for wiring and drilling, and the circuitry is rerouted for your next board.

When it comes to moving a project from paper to prototype to production, nobody does it faster or more economically than Multiwire. To learn more, return the coupon, or circle the reader service number.

MULTIWIRE
10 Andrews Rd., Hicksville, NY 11801

☐ Please send me your 8-page sales brochure.
☐ Please have a salesman call.

Name

Title

Company

Address

City/State/Zip

Telephone

CD

MULTIWIRE® DIVISION
KOLLMORGEN CORPORATION

MULTIWIRE/NEW YORK-31 Sea Cliff Ave., Glen Cove, NY 11542 (516) 448-1428; MULTIWIRE/NEW ENGLAND-41 Simon St., Nashua, NH 03060 (603) 889-0083
MULTIWIRE/WEST 3901 East La Palma Ave., Anaheim, California 92807 (714) 632-7770; MULTIWIRE/ADVANCED MANUFACTURING GROUP 10 Andrews Rd., Hicksville, NY 11801 (516) 938-2000
Multiwire is a U.S. registered trademark of the Kollmorgen Corporation.

CIRCLE 90

updates and/or hardware additions and deletions while it processes application requests.

The operating system supports this procedure through several aspects of update ability. The first component of this ability to update deals with software and data file changes in Unix RTR. Software update can install new or replacement system programs or files, inform the system about them, logically connect them into the system, exercise them in that state, and then commit to or back out of them. Software update is intended primarily for installing fixes or small features that do not disturb the system's architecture.

The second component of update ability, system update, allows program and data changes of a much greater magnitude—up to a complete system replacement. A bootstrap is required to install the changes for any system update. By taking one of a duplicated pair of disks offline, system update can prepare a new release on the offline disk and then reboot with it. Any failures detected during the boot and soak period will automatically switch the system back to the earlier release.

The third component of update ability, hardware, provides the ability to add or remove hardware and related software components to a running system. Growth extends from physically connecting new equipment, such as disk controllers, through informing the system of the availability of the new equipment, diagnosing it, logically connecting it into the system's configuration, and committing it to system use.

Software and system update

Unix RTR is a process-oriented operating system. Processes can be continuously running (ie, they cannot be terminated), or may be stopped and restarted (ie, terminated). Processes that cannot be terminated are the kernel, drivers, and critical application processes. Most Unix processes are processes that can be terminated. Processes that cannot be terminated are updated using function replacements. These allow C language functions or assembly routines to be replaced in the main memory (see Panel, "Software function replacement"). The disk image is left unchanged.

Processes that can be terminated are updated by file replacements which update an executable image on the disk. In both cases it is possible to restore the programs or files to their original state upon operator request (or automatically when a system initialization occurs). The changes can also be made permanent after adequate soaking.

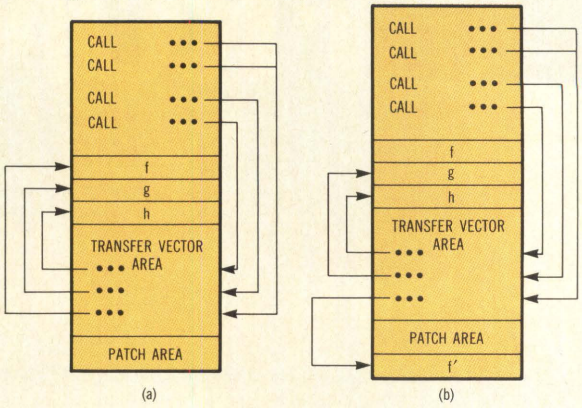
When software development is in progress, the developer must supply only the new source files associated with the update. The software generation system compiles the new source, compares the new object code with the old, determines what kind of replacement to use, and generates the replacement files for the field.

A software update depends on many operating system services and capabilities. One of these calls for primitives to find and write into a process' main store image or change a file temporarily. Another is the support of transfer vectors—an address table maintained on a per-process basis with one address for each function.

System update is a procedure used to introduce an entirely new release into an operational site. Although software update can be made functional in a system without a bootstrap, a system update requires one. In addition, software update is used relatively often to introduce fixes, but system update is used only when a new release or major subrelease is issued.

Software function replacement

Since a process is a collection of functions, the C language function was chosen as the update unit. To avoid making changes to all references in a changed function, all function references are routed through a single reference point called a transfer vector. The Fig (a) is an example of a simplified process image showing the use of transfer vectors. The transfer vector area contains a list of the addresses of the process functions. When a change is made to function f, the new version f' is written into a special "patch" area provided with the process. The corresponding address in the transfer vector area is switched to point f' [see the Fig (b)].



This implementation allows the fix to be backed-out by changing the address in the transfer vector back to its original value. When the fix has been tested and is ready to be applied permanently, the space occupied by f can be reclaimed for future fixes. In addition, decision functions are used with transfer vectors to coordinate the introduction of changes to a set of related functions.

Decision functions allow a process to select the old or the new function based on the signature of the process and the signatures associated with the new and the old functions. With these capabilities, the impact of introducing a new or changed function is restricted to a small, well-defined area of the process, thus making this activity inherently more reliable.



ABLE's ATTACH customers enjoy their spaghetti in the dining room, not the computer room.

ABLE's ATTACH, the breakthrough multi-host terminal switching system for DEC UNIBUS computers that eliminates the spaghetti-like mess of cables in your computer room.

One ATTACH host board and a *single* cable replaces 16 DEC interfaces, and their associated "spaghetti." And it still supports up to 128 terminals on your system without the endless tangle of cables that tie-up your computer room.

One ATTACH host board does the work of many multiplexers. The immediate advantage is a dramatic reduction in mounting space and expansion cabinetry, resulting in significant cost savings. And ATTACH can be located up to a kilometer away from your computer room.

Cook up a system to meet your present data communication requirements with ATTACH. At the touch of a keyboard, terminals can be dynamically *switched* among any combination of VAX and PDP-11 UNIBUS systems. And, ATTACH is compatible with RSX, RSTS/E, VMS and UNIX operating systems.

Expanding your capabilities, or adding terminals is easy with ATTACH. As your requirements grow, simply add modular ATTACH units to your system.

Whether you have 28, or 128 or more terminals, ATTACH has the right recipe for cost-effective connectivity.

ABLE Computer's ATTACH is the most efficient terminal interconnection system on the table.

Contact the ABLE representative near you, or call ABLE toll-free at 800-332-2253.

ABLE COMPUTER

The communication specialists

1732 Reynolds Avenue, Irvine, California 92714. In the Irvine area: (714) 979-7030. Or, TWX: 910-595-1729.

DEC, PDP, RSTS, RSX, UNIBUS, VAX and VMS are trademarks of Digital Equipment Corporation. UNIX is a trademark of Bell Laboratories. ATTACH is patent pending.

CIRCLE 91

In a system update, application-provided code can be introduced to allow application data to be preserved across a bootstrap (to minimize its impact). During a system update, one of the duplicated system disks is taken offline, reloaded with the new release from tape, and the system is then rebooted with the updated disk. Until the second disk is updated with the new release, the system can be rolled back to the old software via a bootstrap.

Hardware update is the ability to add, remove, or modify hardware units together with their related software without incurring system downtime. Hardware update is based on the duplex nature of the hardware units. One unit of a duplicated pair of hardware can be taken out of service, powered down, replaced, diagnosed, and brought back into service again.

Hardware update is assisted by a set of installer and operator procedures. Hardware configuration changes are recorded into the system data base as they occur. Inconsistent hardware configuration changes are detected and disallowed by the data base.

Operator interfaces

Although not specifically related to real time or reliability, the operator interface to Unix RTR is a significant aspect of the system. Its human-factors design allows the operator to make accurate, split-second decisions about system configuration in order to keep a system up, while at the same time minimizing so-called "cockpit" errors. For example, a color graphics terminal is used to display both hardware configuration and software status information in easily digested pictorial form (Fig 3). Similar display screens can be brought up to show the status of critical system resources, disk activity, or application-specific information. These screens not only show status information, but also provide menus that the operator can select from to remove from service, diagnose, and restore to service, any hardware unit.

These displays and menus are maintained by a software system called the display administration process (DAP). In recovery situations where software may not be reliable, a microprocessor-based configuration controller gives the operator instantaneous and complete control over hardware configuration. This controller has its own menu page as well (Fig 4).

Human-factors design allows operators to make accurate, split-second decisions about system configuration, while it minimizes "cockpit" errors.

To gain a better understanding of various fault detection and recovery approaches, it is helpful to contrast the approaches taken by Unix RTR and 3B20D with those taken by other vendors. Of course, error detection and recovery through the use of redundant hardware and software components are central to the design of a highly reliable and available system. It is important that such systems detect hardware and software faults as soon as possible and then take appropriate recovery actions depending on the nature of the faults.

Hardware errors can be broadly classified into storage, transmission, and computational errors. The use of error-correcting codes (eg, Hamming code) to protect against memory faults, as well as the use of parity bits in data transmissions, are

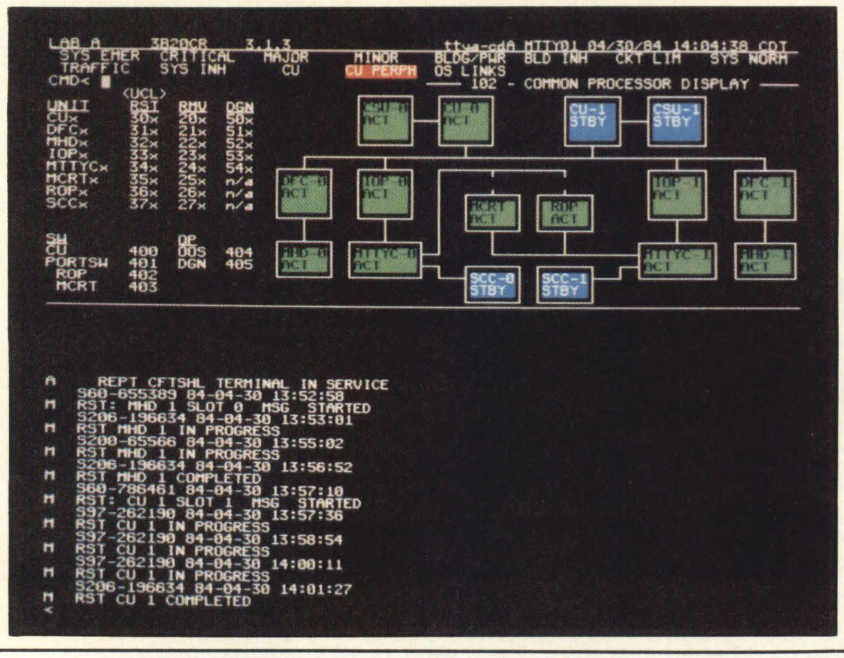


Fig 3 The 3B20D teams with a color display terminal to provide hardware and software status information. A typical hardware status report is shown. Similar screens may show other critical system resources that aid in monitoring computer reliability and operation.

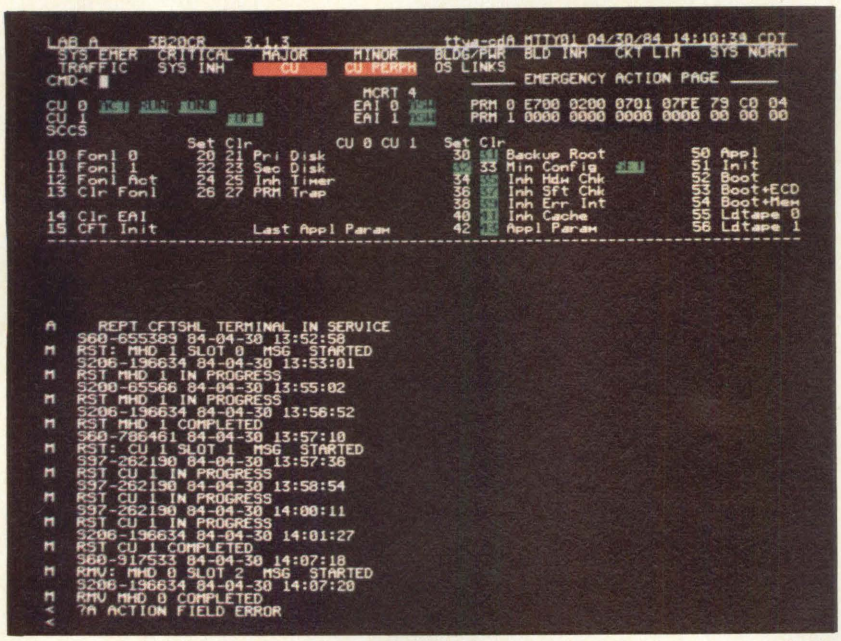


Fig 4 Under certain circumstances, a microprocessor-based configuration controller provides the 3B20D operator with information about, and control over, the machine's hardware configuration. This controller has its own menu page, which may be displayed on a color monitor.

common among all fault-tolerant systems. The detection of computational errors, (eg, errors in adding two numbers), is much more complicated and is usually achieved by voting schemes or a multiprocessor configuration with duplicated ALU.

With voting schemes, three or more processors work on a task and the system uses the result that gets a majority of votes. A single failure in the voter will bring down the system. In a multiprocessor configuration with duplicated ALUs, the approach taken by Unix RTR, a processor is declared to be out of service when the results of the duplicated ALU disagree. When this occurs, processes running on the faulty processor must be restored on one of the remaining processors. In other systems, this is achieved by creating process pairs and saving process states periodically.

The architecture allows processes to handle their own software faults.

The primary process is the active program. The backup process is a passive copy of the primary process and resides in a different processor. Before performing any critical function, the state of the process is saved or checkpointed. The saved information includes data and processor status information that the backup process needs to complete the function should a failure occur.

With this design approach, fault tolerance is not software-transparent and cooperation from the application software and system software is required. In Unix RTR, a bidirectional memory update capability allows process states to be checkpointed continuously. Thus, when a hardware fault occurs, the surviving processor can easily take over all the processes in the system without user-provided checkpointing.

Since software faults are permanent faults (ie, the fault is triggered whenever the same set of conditions are presented), most fault-tolerant systems do not handle software faults adequately. Since it is virtually impossible to eliminate all errors from any nontrivial software, a system with high availability must be able to recover from software faults gracefully. The architecture of Unix RTR allows processes to handle their own software faults (eg, protection violation, out of range addresses, etc) through their own fault entries. System processes that are transaction-oriented support a rollback and recovery strategy to deal with such faults. If a rollback is insufficient to clear the problem, progressive initializations will be automatically invoked to clear the problem.

The unique capabilities of Unix RTR include audits and overload control which further enhance its availability. Audits verify the validity of critical system data and retrieve lost system resources. Overload control ensures that system resources are available to critical applications and that the realtime performance does not fall below a predetermined limit, even under heavy load situations.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 725 *Average 726* *Low 727*


NCR
Disk
Communications
Fault
Battery
On
Off

FOR OEMs, THE GREATEST STRENGTH OF OUR TOWER IS OUR TOWERING STRENGTH.

To build a successful system, OEMs need more than just a strong product. You need a strong product backed by a strong company.

A company like NCR.

NCR designed and manufactures Tower 1632 especially for OEMs. But that's only the first of many reasons why you should put your name on our famous shape.

NCR is a high volume, high speed, high reliability manufacturer. We're deeply committed to 16 and 32-bit VLSI technology. With truly significant R&D expenditures that only a multibillion-dollar, international corporation can sustain. And long-term dedication to the UNIX* market.

That means you can rely on getting all the product you need from us, when and where you need it. It means that you get quality control from a company that knows real quality and can afford to build it into every product. For example, our field engineers participate in the design of our products to ensure exceptionally cost-effective serviceability.

Our award-winning design engineers packed Tower's 7″-wide cabinet with up to 2MB of ECC memory, 7 controller slots, standard interfaces, power-fail recovery and 92MB of Winchester disk storage. Not to mention maximum software compatability and flexibility.

But the real reason Tower is becoming the industry standard for multi-user, 16-bit and 32-bit UNIX-based systems is the strength of NCR.

Which proves that the shape you're in is determined by the company you keep.

*UNIX is a trademark of Bell Laboratories.



BUILT FOR SYSTEMS BUILDERS.
TOWER 1632.

NCR
OEM Systems Division

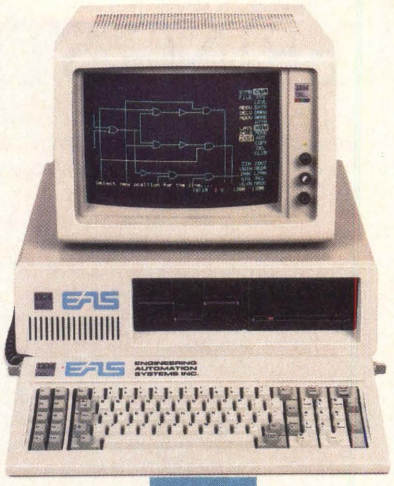
NCR Corporation, World Headquarters, OEM Systems Division, Dayton, Ohio 45479. Nationwide (800) 222-1235. In Ohio (513) 445-2380. In Canada (800) 268-3966.

CIRCLE 92

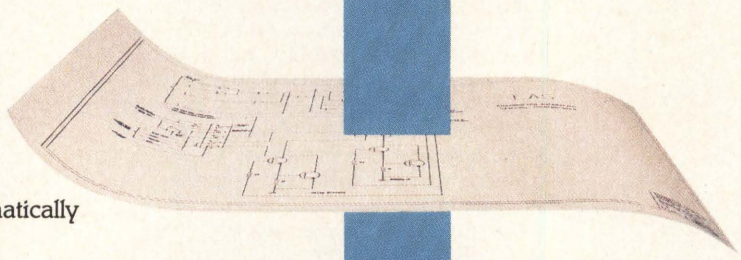
Think Design Capture

No more drudgery of designing circuits or drawing schematics by hand. Or tying up an expensive workstation just for CAE tasks. Hand the job over to EAS . . . to maximize your productivity through integrated CAD/CAE.

Design circuits on the EAS/300 with PC-CAPS™ — a true hierarchical design system. Create a database that constantly keeps track of components and connectivities . . . with real-time continuity checking and instant wire trace capability. Editing and viewing is easy. Verify circuit performance with the PC-LOGS™ interactive, event-driven logic simulator.



Download net-list data through ETHERNET® or via RS-232 interface to an EAS/770 CAD workstation for complete PCB layout. Back-annotate from the EAS/770 to the EAS/300. Automatically draft schematics on your plotter.



Think PCB Layout

With the EAS/770 and its PCX™ software, design complex circuit boards with automatic/interactive placement and routing, and real-time design rule checking. Get outputs for manufacturing. Transfer design data throughout a wide network of EAS/770 and EAS/300 workstations in a true integrated CAD/CAE environment.

The EAS/770 — a responsive, menu-driven XENIX® (UNIX™) based PCB CAD system. The EAS/300 — a personal engineering workstation for desktop CAE.



THINK EAS.
THE TOTAL SOLUTION.

EAS ENGINEERING AUTOMATION SYSTEMS INC.

23 Industrial Park Road Middletown, CT 06457 Phone: (203) 632-0080

PC-CAPS and PC-LOGS are trademarks of Personal CAD Systems, Inc.

ETHERNET is a registered trademark of XEROX Corporation.

PCX is a trademark of Engineering Automation Systems, Inc.

XENIX is a registered trademark of Microsoft Corporation, and is derived from UNIX System III under license from AT&T.

UNIX is a trademark of Bell Laboratories.

CIRCLE 93

FAULT-TOLERANT MINI NEEDS ENHANCED OPERATING SYSTEM

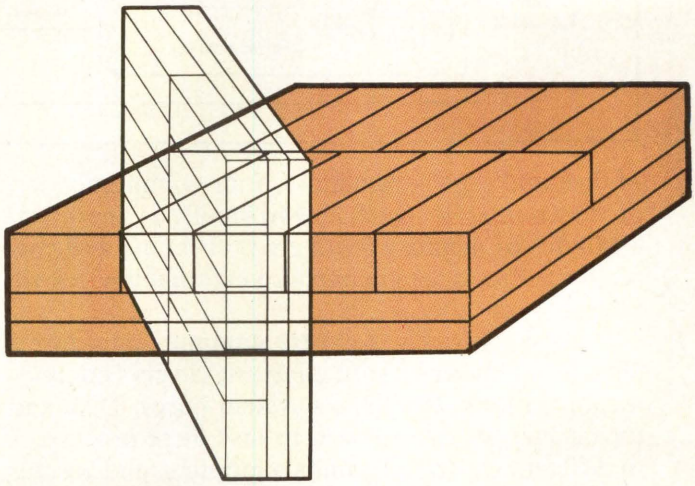
Unix System III application programs run in a user-transparent, expandable minicomputer environment.

by Samuel D. Glazer

Minicomputer system designers often require a fault-tolerant architecture to make their systems suitable for transaction processing or other environments where downtime is unacceptable. One approach calls for multimicroprocessor-based, distributed hardware to run an enhanced version of the Unix System III operating system. With a properly designed operating system of this type, a variety of application software can run transparently without the need for special user-created, fault-tolerant code.

One operating system using software that is designed to work well with a fault-tolerant hardware architecture is called Auros. Derived from Unix, this operating system is part of the Auragen System 4000 fault-tolerant minicomputer. To handle fault tolerance, the Auros kernel is set up to control the message-passing communications that allow an active hardware cluster to take over for a failed one.

The Auragen System 4000 minicomputer comprises 2 to 32 tightly coupled, bus-connected multiprocessor units (clusters). A bus-connected distributed architecture provides modular hardware growth as well as fault tolerance. The multiplicity of dedicated microprocessors ensures that each computing microprocessor is off-loaded from overhead chores.



Each Motorola 68000 family-based cluster contains an executive processor module (to run the Auros operating system for interprocess and intercluster communications), a work processor module (to execute application programs), and a 2-Mbyte memory module (see Fig 1). Clusters expand to contain four disk/tape processor modules (controlled by a bit-slice microprocessor for high speed handling

Samuel D. Glazer is director of software development at Auragen Systems Corp, Two Executive Dr, Fort Lee, NJ 07024. He holds a BS in mathematics from Columbia University.

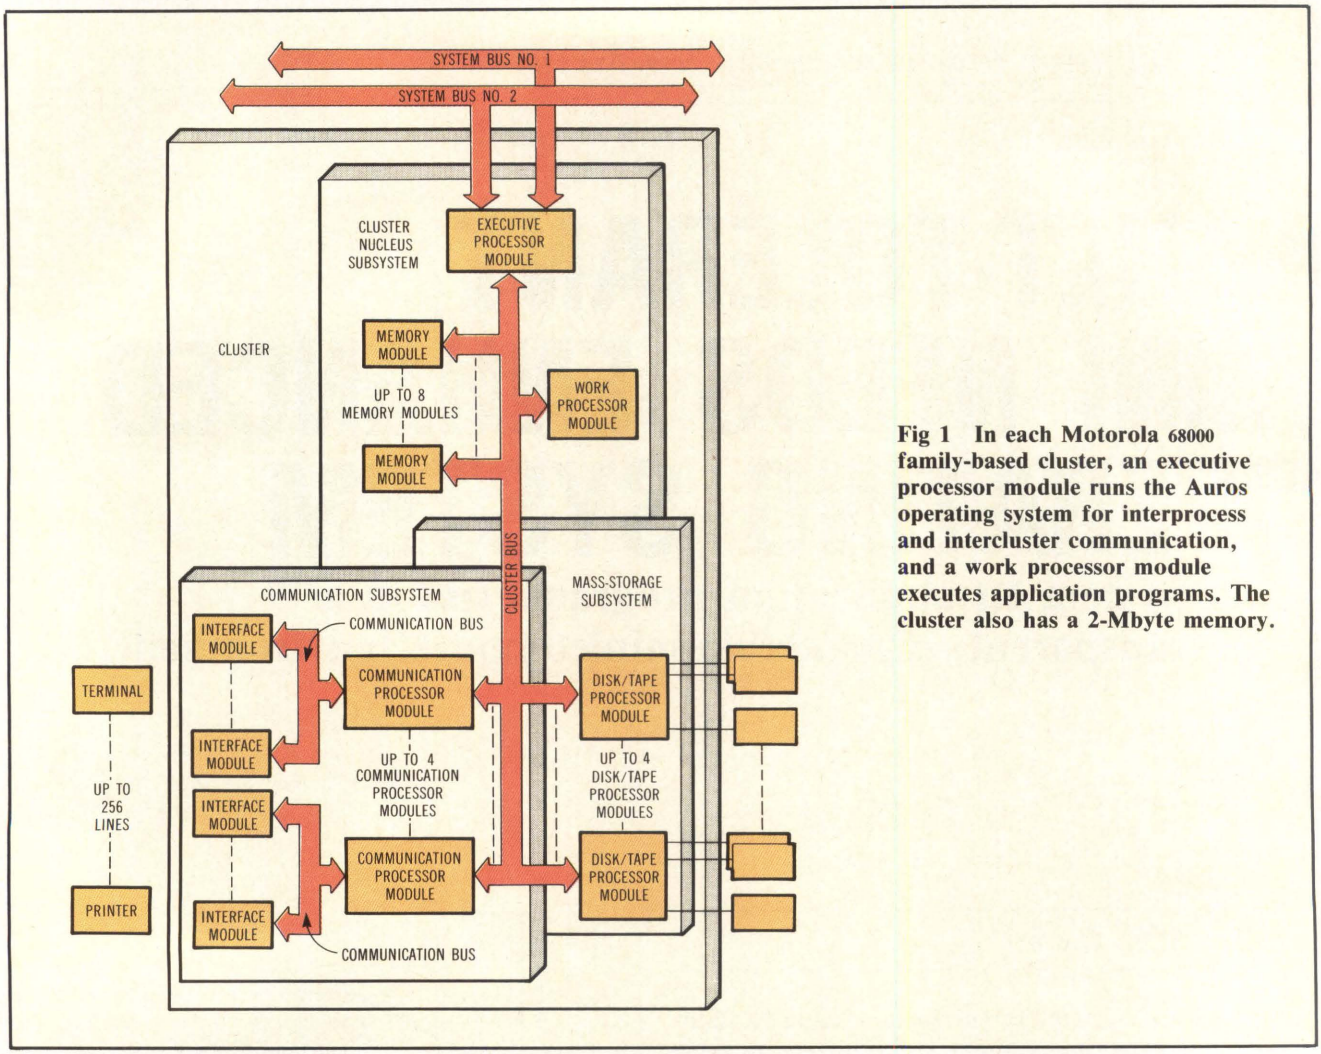


Fig 1 In each Motorola 68000 family-based cluster, an executive processor module runs the Auros operating system for interprocess and intercluster communication, and a work processor module executes application programs. The cluster also has a 2-Mbyte memory.

of disk and tape commands), four communication processor modules (each able to handle 64 lines), and 8 Mbytes of memory (with error checking and correction). Processors connect to external devices via interface modules.

Processors within a cluster communicate over a 20-Mbyte/cluster bus. In contrast, clusters communicate via two 16-Mbyte/s system buses. Disk and tape drives are dual-ported to disk/tape processors in different clusters. Terminals, printers, and various other I/O devices connect to communication processors in different clusters via 4-Mbyte/s communication buses.

A fault-tolerant computer needs more than just the hardware fault tolerance of distributed clusters, however. Its software must enable the clusters to handle the fundamental job of a fault-tolerant computer. In this job, backup software executing in a backup cluster takes over for the failed software process (executing in a cluster whose hardware fails). Backup hardware and its software must replace failed hardware without losing data or program steps, and without effecting an application programmer or end user. The user should be unaware of when failures occur, which cluster executes the application, or when files may be active in another cluster dedicated, for example, to file service (for efficiency).

All fault-tolerant computer systems require hardware and software resource duplication. They differ, however, in the degree to which they allow the backup resources to perform useful work when there is no fault; in the performance loss of the primary computer in maintaining the backup; in recovery time after a fault; and in the ease of programming fault-tolerant applications. They are similar in that they all require that the backup for the primary process (whether or not it executes while the primary process is going on) be able to continue without data loss after a failure (see Panel, "Compare and contrast"). For Auragen's System 4000, Auros is the keystone for addressing these considerations.

Kernel control

The kernel-based message system handles all Auros interprocess communication for normal hardware system operation and fault tolerance. For example, the kernel handles message passing between

UNFORGETTABLE



At last, random access memory that won't let you down. Introducing Mostek's MK48Z02. We call it ZEROPOWER™ RAM. You'll call it unforgettable.

Because that's exactly what this new non-volatile RAM does. It keeps its memory. Thanks to the integration of advanced CMOS technology, ZEROPOWER RAM has an on-board, long-lived lithium energy source. That means ZEROPOWER RAM does not forget, even if the power goes down. ZEROPOWER RAM can keep on remembering for over 10 years without power, drawing less than one nanoamp at room temperature in the power-down mode. Plus it can be reprogrammed quickly and efficiently, as often as you want, without the wear-out mechanisms associated with E^2PROMs.

Also, this amazing new RAM can be read to and written from like the static RAM it is. Data access is 150 nanoseconds with an active power consumption of 330 milliwatts. And the MK48Z02 can replace existing 2K x 8 static RAMs, directly conforming to the popular bytewide 24-pin DIP package. No other components are needed.

So when you are looking for memory power that lasts, remember Mostek's MK48Z02 ZEROPOWER RAM. It's the unforgettable memory.

For more information, contact Mostek, 1215 W. Crosby Road, MS2205, Carrollton, TX 75006, (214) 466-6000. In Europe, (32) 02/762.18.80. In Japan, 03/496-4221. In the Far East (Hong Kong), 5-681157.

ZEROPOWER is a trademark of Mostek Corporation.

UNITED TECHNOLOGIES MOSTEK

CIRCLE 94

primary processes executing in a distributed environment and also ensures that the secondary process has all the information it needs to bring itself fully up-to-date if failure occurs.

Unix System III does not have a kernel-based message-passing system. Since it was not designed for a distributed, multimicroprocessor architecture, it does not have the separate system servers necessary for an efficient distributed processing environment. In addition, its file system is not reliable enough for transaction processing applications since it can lose files and become inconsistent, and it does not efficiently manage the processes needed (eg, for transaction processing).

Auros handles these deficiencies in a manner fully compatible with Unix System III software. Each of the System 4000 clusters runs its own cluster-resident copy of the Auros kernel, which contains a so-called "queue-and-count" message system. The message system provides interprocess communication between primary processes and supports fault tolerance by ensuring that every cluster-resident executing program or primary process has a fully informed backup copy on a different cluster. To do its job, Auros supports as many as 256 concurrent tasks in each cluster. In addition to message passing, it features new system calls for transaction processing environments, faster process creation and process switching, and demand-paged virtual memory.

Choice of two

There are two types of Auros software processes. User processes communicate all I/O information by message passing. They have no direct access to any computer peripheral equipment. Actual device I/O (in systems with peripherals) is handled by requests sent to the peripheral server processes. Peripheral server processes take care of specific devices that they access via system calls.

Fault tolerance for user processes is based on the concept that if two processes start out in an identical state and receive identical inputs, they must perform identically and produce identical outputs. As previously stated, each primary software process has an inactive backup resident in a different cluster. A backup process is kept nearly up-to-date, and is provided with the information needed to bring itself fully up-to-date with the primary state and to continue execution as the new primary should a failure occur. Thus, a primary and its backup are initially identical, all the input messages to the primary are available to the secondary, and the backup can recompute to catch up to where a primary failed (by using the same messages the primary used).

Complete recomputation by the backup is not necessary since the primary and backup processes are periodically resynchronized. Between synchronizations, when the backup and the primary are different, all messages to the primary are made available to the backup. Upon synchronization, all messages previously read by the primary process are discarded by the secondary one. If the primary fails, the backup executes—rolls forward—from the last synchronization point (using the saved input).

The message system supporting the relationships between the primary and the secondary fault-tolerant equipment is embedded in the kernel. This message system provides and controls interprocess communication, initiates the creation and deletion of backup processes, and controls the periodic synchronization

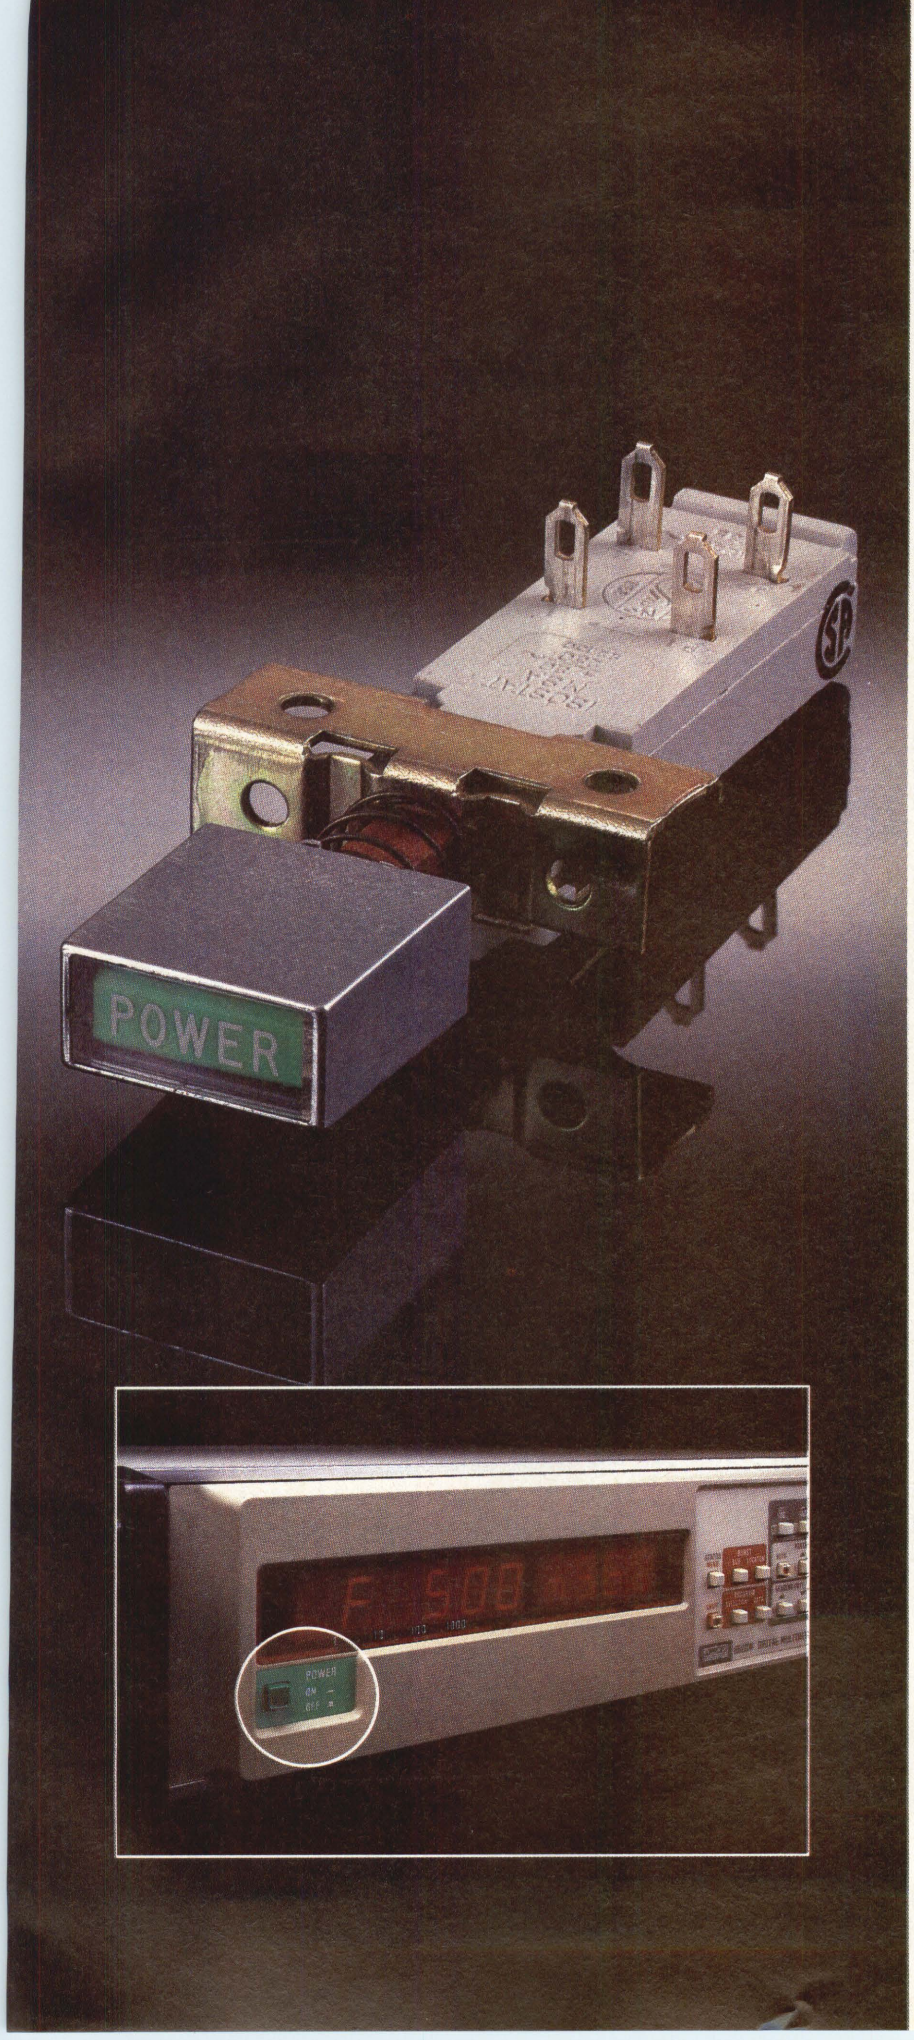
Compare and contrast

There are two approaches to backing up primary processes that are related but slightly different. One approach requires that the primary process and its backup execute simultaneously on two independent CPUs. If one CPU fails, the other continues without interruption. The duplicate hardware provides no computational advantage, there are no special fault tolerance programming considerations, recovery time is instantaneous, and the performance of the primary is not taxed by the secondary's presence.

An alternative is to maintain an inactive or nonexecuting backup process that executes on a secondary processor upon failure of one processor. The secondary can do useful work while it backs up the primary. There are differences in how the backup is maintained, with consequent differences in primary and backup resource use and recovery time. For example, the state of the backup, as represented by its data space, must either be identical to that of the primary or be capable of being made so (the Auragen approach).

The data space may be kept identical by means of "checkpoint" software which copies the primary's data space to the secondary whenever the dataspace changes. In contrast, Auragen has opted to provide the means for updating the data space by its "queue-and-count" interprocess message system software. When there is a cluster failure, the backup process automatically recreates a data space identical to the primary's at the time of the crash, so a backup process can continue processing as if no failure occurred.

Queue-and-count makes available to the backups the same messages it delivers to the primary. This queue of messages is used during recovery by the backup to recreate the primary data space of the process. In contrast, checkpoints are large messages sent from the primary to the backup each time a normal interprocess message is sent or received. Since user programs must supply checkpoints, checkpoint instructions must be embedded in user programs. Application-program transparent queue-and-count software requires no special programming.



ITT Schadow's Power Switches

Multinational approval and high surge ratings are good reasons to use ITT Schadow Power switches in your computers and peripherals.

If you market your equipment overseas – and these days, who doesn't? – ITT Schadow's Pushbutton switches are worth specifying for their multinational approvals alone. And even if you never sell a piece of hardware beyond the borders of the U.S.A., you are bound to appreciate the ruggedness and high performance of these outstanding switches.

ITT Schadow's MSA has a dual rating of 10Amps at 125VAC and 4Amps at 250VAC, with an 80Amp surge in either instance. (Switches with a 64Amp surge rating are also available). A shutdown option lets a central location shut off terminals by remote control.

For details on ITT Schadow Power Switches – or for information on ITT Schadow's custom switchmaking and worldwide sourcing – mail the coupon below. Or contact your local Authorized ITT Schadow Switch Center.

FREE SAMPLES

☐ I buy or specify switches. Please have a rep call me about a free Power switch Schadow Box.

☐ I am not ready to talk with a rep yet, but please send a full-line switch catalog.

CD-8

Name _____

Title _____

Company _____

Address _____

City _____

State _____ Zip _____

Telephone ( ) _____

MAIL TO: ITT Schadow Inc., a subsidiary of ITT Corporation, 8081 Wallace Road, Eden Prairie, MN 55344
(612) 934-4400

Schadow ITT

Circle 95 for Catalog
Circle 138 for Sales Engineer Call/Samples

of primary and backup. It further guarantees that the backup can take over in case of failure if it has all the information it needs, and that the backup will interact correctly with the system.

The message system ensures that during normal process execution, all messages sent to the primary (which were unread or arrived after the last synchronization) are available to the backup. It also insists that the primary's state as of the last synchronization time is accessible to the kernel (which controls the backup's processing cluster), and that during a recovery, the backup process reads the available messages in exactly the same order as did the primary. Finally, the message system must see that during recovery the backup will not send any messages already sent by the primary.

Each message sent from one primary process to another is actually sent to three destinations. These are the requested primary destination process, its backup, and the sending process backup (Fig 2). If a message is sent to more than one location, either all locations or none must receive it. The software must guarantee that when a message arrives at each of its three destinations, it is never interleaved with that of any other message. This rule means that a primary and its backup always receive messages in the same order. In addition, if two messages are sent, the first message must reach all of its destinations before the second message arrives at any of its destinations.

A primary process and its backup will always receive messages in the same order.

Each of the three destinations has software that makes use of each received message in a different way. For example, at the primary destination the message is queued for reading by the primary destination process. In contrast, at the backup destination the message is queued and saved for the destination process backup. It is to be read only upon roll forward after a failure. At the sender's backup, a count of the messages sent since the last synchronization is incremented and the message is discarded. Thus, every backup process has a queue of the messages sent to its primary and a count of the messages sent by its primary. The primary keeps a count of the number of messages it has read since the last synchronization.

Stand up and be counted

How software control is transferred from one program to its backup when a fault occurs is important because the method determines how much useful work the backup can do when there are no faults to handle. The control transfer technique also

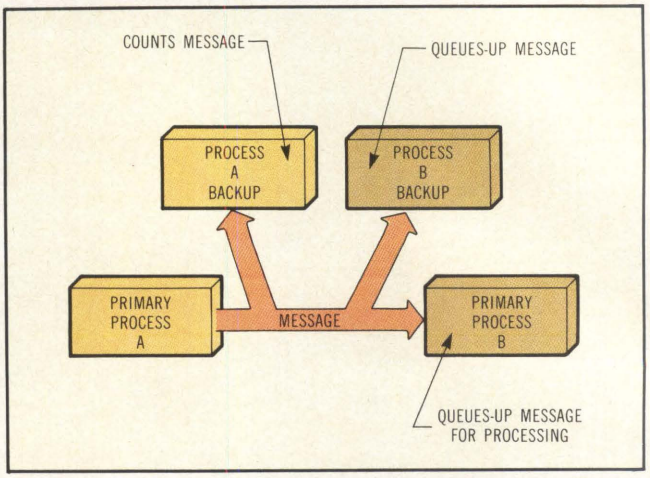


Fig 2 When a message is sent by a primary process to its next process, the primary's backup receives a copy, as does the backup for the next process. If one message fails to arrive, the system bus ensures that none arrive, so the queue-and-count software does not fail.

determines how fast the fault-tolerant computer recovers from a crash.

Auragen was not satisfied with checkpointing for transferring software control (as explained in the Panel, "Compare and contrast"), and developed the queue-and-count design to manage System 4000 interprocess communications. As a manager, queue-and-count performs two chores: ensuring that the 4000's distributed processor clusters appear to the system user to be a single system; and ensuring that secondary (backup) programs always have the correct, latest information so they can take over should a hardware failure that takes out a cluster occur. In the queue-and-count design, whenever a primary process receives a message from Auros, a corresponding backup process also receives the message and saves it in its "queue"—hence the first part of the queue-and-count name. It remains in the queue until the secondary must take over from the primary, at which time the message is processed. Since the secondary must also know what the primary has accomplished up to the time when a failure occurs, it also keeps a count of the primary's outgoing messages (sent to the next primary process). With this "count" part of the queue-and-count scheme, the secondary can avoid sending duplicate messages.

A fault-tolerant system with three primary and three secondary processes illustrates the queue-and-count concept. This example illustrates how a terminal server process handling I/O, an application program, and a file server that handles reads and writes, normally performs and recovers from failure (see Panel, "How queue-and-count software handles recovery"). Remember that the key to recovery is to reproduce the state of a primary at the moment it failed and then continue with normal operations. Except for a several second pause during recovery, the user does not know the fault-tolerant system has

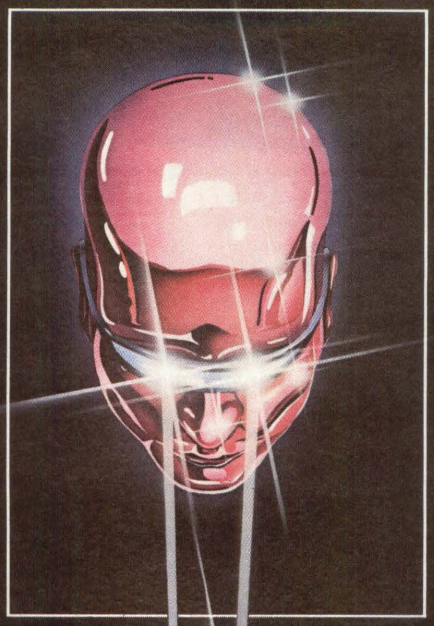
The Convergence Factor.

Convergence: the single most critical factor in color CRT performance.

Until now, Delta-gun tubes were the best way to achieve near perfect convergence, but only with costly adjustment electronics. Meanwhile, many in-line tubes are plagued by perceptible misconvergence. Which can lead to poor picture quality. A poor quality image for your product. And poor, bleary-eyed operators.

The Panasonic achievement: low cost in-line color CRTs with better-than-Delta convergence performance.

Without complex adjustment electronics . . . and none of the convergence drift inherent in active correction systems. At last, high resolution in-line tubes with stable performance that stands up to the ravages of time and tough office/industrial environments.



How did we do it? With a preconverged in-line tube/yoke combination unlike any other. Our precision S/ST (saddle/saddle toroidal) deflection yoke is ideally matched to each tube, for near perfect convergence, high repeatability and stability over a wide range of operating conditions.

We combine it with a specially-designed OLF (overlapping field lens) gun and unitized grid construction, providing spot uniformity across the entire screen and near-Delta resolution.

The result: a triumph over the convergence factor. Find out what it can do for your next color terminal or monitor, and ask about our full line of quality color and monochrome CRTs. Write or call: Panasonic Industrial Company, Electronic Components Division, One Panasonic Way, Secaucus, N.J. 07094; (201) 348-5278.

The achievement of Panasonic high resolution in-line color CRTs.

Panasonic
Industrial Company

CIRCLE 96

How queue-and-count software handles recovery

Suppose a primary terminal server sends a message (M1) to a primary application program which, in turn, sends a message (M2) to a primary file server to write a record for a disk (see Table, step 1). In the queue-and-count procedure, the message system sets the terminal-server backup count indicator to 1 (indicating that the primary terminal server has sent one message). The message system stores message M1 in the backup application program's queue and sets its count to 1 since the primary application program has sent a message to the primary file server. For the file-server backup, the message system queues M2 but sets no count indicator because the primary file server sent no messages. The primary terminal server then sends message M3 to the primary application program and operations occur as shown in step 2 of the Table.

The primary terminal server (step 3) sends an M5 message which is received by the primary application program. But, assume that the primary application program fails to send M6 to the primary file server. The terminal-server backup has a count of 3 (meaning that the primary terminal server has successfully sent three messages). However, the count on the secondary application program is 2 (since the third message was not sent). The application program backup has M5 queued along with M3 and M1. The secondary file server is unchanged since no step 3 message has been received.

Recovery starts when the previously inactive application program backup becomes active and executes like any other process after checking its queue. Messages in the backup are read in the same order as they were received in the primary application program. This reading order and the count availability ensure that redundant messages are not sent and that the next message to the backup file server is M6.

To initiate recovery, the application program backup (with a count of 2) reads message M1 in its queue, executes it, and generates the appropriate message to be sent to the file-server backup. This new M2 message is not sent to the file-server backup because it is a duplicate of a message the file-server backup already has and is recognized as such because the application program backup's count is 2. Queue-and-count requires a count of 0 before a new message is transmitted since 0 guarantees no previous message transmission.

Since M2 is a duplicate, queue-and-count discards it. The counter on the application program backup is decreased by one and message M3 (the second message received by the primary application program) is read by the secondary application program. M3 is read, executed, and used to generate M4, but M4 is not sent since the count is 1.

The backup application program count is then decreased to 0, the program reads the next message (M5) in its queue, executes it, and generates message M6. Because the application program secondary's count of messages sent by the application program primary to the file server primary is 0, the queue-and-count system knows that the primary file server never got M6. So, the message system sends M6 to the primary file server, which writes its record to a disk, and fault recovery is complete.

Queue-and-Count Example

Queue-and-Count Step	Terminal Server	Application Program	File Server	Terminal Server Backup	Application Program Backup	File Server Backup
1	Sends M1	Receives M1 Sends M2	Writes to File	C = 1	C = 1 Queues M1	Queues M2
2	Sends M3	Receives M3 Sends M4	Writes to File	C = 2 C = 2	C = 2 Queues M1, M3	Queues M2, M4
3	Sends M5	Receives M5 Does not send M6	—	C = 3	C = 2 Queues M1, M3, M5	Queues M2, M4
4	—	—	—	—	C = 2 Read M1 Generate M2' Delete M2' C = 1 Read M3 Generate M4' Delete M4' C = 0 Read M5 Generate M6 Send M6	Queues M2, M4 Receives M6

Legend: M = Message
C = Count

Note: Start at queue-and-count step 1.
Read Horizontally left to right.
Go to next step. Repeat until end.

failed. Of course, the flawed hardware signals its disability to system maintainers.

When a primary process fails, the secondary process uses messages to retrack the primary one's actions until it is in the same state as the primary process before it failed. The secondary process then continues with the calculations. If a failure occurs after many messages have been transferred, it might seem that device restoration to its original state is inefficient. It might also appear that an unbounded message queue wastes memory resources. This is not true, however, because both primaries and the secondaries are periodically synchronized. After each resynchronization the queue restarts, as does the count.

Save that state

A primary and its backup are automatically synchronized whenever the primary has read more than a system-defined number of messages. Synchronization is also implemented if the primary has executed for more than a system-defined amount of time since the last synchronization. Any changes in the address space of the primary (since the last synchronization) are stored so that they are available to the backup. This process is made possible by cooperation between the message system and the computer system's paging mechanism.

Processor synchronization must deal with both real and virtual processor memory. The processor's normal page-fault mechanism ensures that pages changed in real memory since the last synchronization are updated in the virtual memory. The primary also synchronizes its virtual memory with that of the backup when the message queue reaches a system-predetermined value. For efficiency, only the virtual memory that has changed since the last synchronization (that has not been updated in the last real memory page file), is sent to the secondary.

Once the memory space changes are stored, a synchronization message containing certain state information is sent directly to the backup processing unit's kernel. This state information includes a count of the number of reads done by the primary since the last synchronization. The count's availability allows any messages saved for the backup, but already read by the primary, to be discarded. After synchronization the backup will have the correct set of messages available. The arrival of a synchronization message also causes the backup's count of the messages sent since the last synchronization by the primary to be zeroed.

Another concern is that the backup not send (to the next process) any messages that were generated by the primary between the last synchronization and a failure. Remember that the third message destination is the sender's backup, where the message is counted and discarded. So every time the backup (which has become the new primary) begins to execute code and send a message, it checks the value of this count. If the count is positive (meaning this message was already sent by the primary), the count is decreased and the message is not sent. In contrast, if the count is zero, the message is sent.

Synchronization has certain advantages over the checkpointing procedure for backup updates. For one, it is automatically initiated by the operating system, making user program instructions unnecessary. For another, synchronization is needed less frequently than checkpointing. Also, primary processes are not slowed by synchronization, as they are by checkpointing, because they need not wait for synchronization to complete. After synchronization, the contents of message queues and counters are set to zero and any failure recovery proceeds from the new synchronization point.

The operating system differs from Unix in several other respects, all of which are designed to enhance fault tolerance. To allow the 4000's distributed hardware to gain the efficiency benefits of distributed client/server software, the operating system is designed with function separation. Unlike the Unix kernel, the kernel only controls local functions such as memory management, resource control, process scheduling, peripheral, and message handling. Backup servers (peripheral and system) handle global resource management for such services as terminal I/O and file and page management (Fig 3). Communication between servers and any user processes also uses the message system. Peripheral servers which are associated with either logical or physical devices, receive messages in the normal way but must be able to execute certain system calls to control their associated devices.

The operating system automatically initiates synchronization, thus making user program instructions unnecessary.

The result of this separation of global and local resources is, for example, the distribution of front-end (terminal I/O) and backend (file I/O) services across multiple microprocessors and clusters. These servers require somewhat different backup and synchronization schemes than user processes.

Peripheral- and system-server backup processes are created when the primary process executes. In contrast, with user processes, backups are created only when necessary to ensure fault tolerance. Peripheral-server synchronization is different from user-server synchronization since peripheral servers must be core-resident rather than paged, and they communicate with devices directly rather than by message.

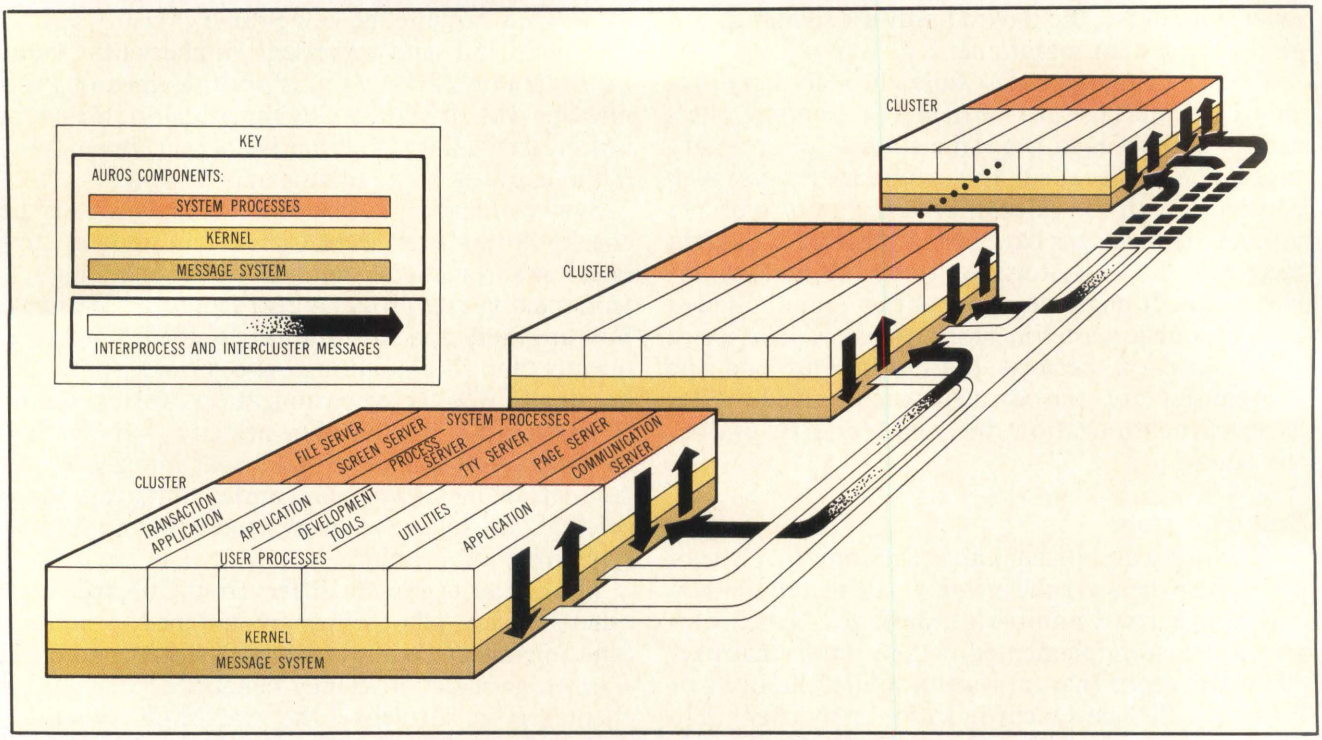


Fig 3 All processes and clusters send and receive messages via the kernel-based message system. Division of the cluster processes into user and system processes allows the kernel to handle local operations while the servers handle global ones.

System-server processes keep track of global system resources via tables in their address space. They are backed up, communicate via message, and execute in the same way as user processes. They differ in that they can be initiated only by the operating system.

All messages transmitted back and forth between processes (whether user or peripheral) are sent by means of a software mechanism known as a channel. Each end of a channel is defined by means of an entry in a cluster's local table known as a routing table. A cluster's routing table resides in kernel space in cluster main memory and is maintained by message system code.

A special process server responsible for system load balancing ensures optimal minicomputer hardware performance. With its proprietary algorithm, the process server determines the cluster on which an application should run. The process server takes into account cluster CPU and memory availability. It also checks to see if the program is already resident in a cluster so that the new process can share resources. The process server, by keeping track of where all processes (primary and secondary) reside, plays a critical part in locating and bringing up secondary processes should a failure occur.

The computer system designer or system integrator opting for a distributed minicomputer is concerned not only with fault tolerance; the easy generation of the application software that must run on the hardware is also a concern. So, among other Unix utilities, Auros supports Programmer's Workbench, file and string manipulation programs, text editors, and document-formatting packages.

In addition to these standard tools, there is a screen manager that allows display creation, a menu system, and a database manager known as Auralate that supports data integrity and concurency control. This relational database system is based on the high level, English-language-like query language SQL. A transaction processing management system (TPMS) to speed the development of transaction processing-intensive applications is also included.

Software development languages provided include C (the Auros language), Cobol (ANSI 74 with an interpreter), Fortran (ANSI 77), Pascal (conforming to the proposed International Standards Organization version and UCSD extensions), and interpreted Basic. All languages include compilers, source debuggers, and runtime libraries, as appropriate.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 728 *Average 729* *Low 730*

First "Graphics Ready" monolithic raster scan DAC.

125 MHz, 8-bit, 800 mW Glitch-free DAC for only $39*

The TDC1018 from TRW is the first monolithic 8-bit, 125 MHz D/A converter that's "graphics ready." This means you can skip the analog buffer amplifiers, video sync and blank inserters or deglitching circuitry—they're on the chip. And because the TDC1018 uses the TRW 1 Micron OMICRON-B™ process, it features higher reliability, improved speed-power product and pixel-perfect signal conversion.

The TRW TDC1018 is your complete solution to raster scan applications, color graphics, digital synthesizers, automated test equipment and high speed, high resolution video system designs.

Our compact design fits into a small 24 pin DIP, and is suitable for your large volume applications. Consider these outstanding features:

- 1.7 ns rise time;<5 ns settling time
- Unique design virtually glitch-free
- On-chip latching for easy interface of blank, sync, bright and reference white
- Differential logic for maximum speed, minimum noise



- Protection circuitry prevents power supply spikes at output
- Directly drives standard 75Ω load
- Multiplying capability allows gain adjustment
- Differential clock reduces power supply ripple effect
- Differential outputs produce true and inverse video
- ECL data and control inputs
- 1/2 LSB linearity

CIRCLE 97

For a full data sheet or immediate information about our full-feature DAC, call or write:

LSI Products Division
TRW Electronic Components Group,
P.O. Box 2472, La Jolla, CA 92038,
619.457.1000

Outside U.S., call or write:
TRW LSI Products
Europe, Konrad-Celtis-Strasse #81,
8000 Munchen 70, W. Germany,
089.7103.0

Kowloon, 3856199; Tokyo, 4615121;
Taipei, 7512062

*U.S. price in 1000s

© TRW Inc. 1983—TRS-3109

TRW

LSI Products Division
TRW Electronic Components Group

The K105-D analyzer with Trace Control™ helps you find faults quickly and easily.

For most engineers, the frustration is not in solving the problem. But in finding it.

Not anymore. Take the Gould K105-D logic analyzer. It can help you isolate errors and their causes quickly and easily. Because now you can capture and examine several possible cause areas with a single pass.

Advanced program tracking and data capture.

Flexibility is the key.

With 8 levels of Trace Control, you can follow every bend and turn in your program flow. No matter how complex. And capture widely-separated slices of code better than any other triggering method.

Although faults and their causes may be separated by thousands of lines of code, the K105-D records only what you want to see.

Once captured, the K105-D brings these windows of information directly to the screen for your analysis.

This selective recording conserves memory. While it helps you get at the problem—and solve it—in less time than ever before.

An analyzer that speaks your language.

The K105-D speaks simple assembly language, not complex object code. Quite fluently, we might add.

With our disassembly modules for the 68000, 8086, 8080, 8085A and Z80® B microprocessors, you need only one analyzer for debugging most popular processors.

Clear, concise pattern definitions.

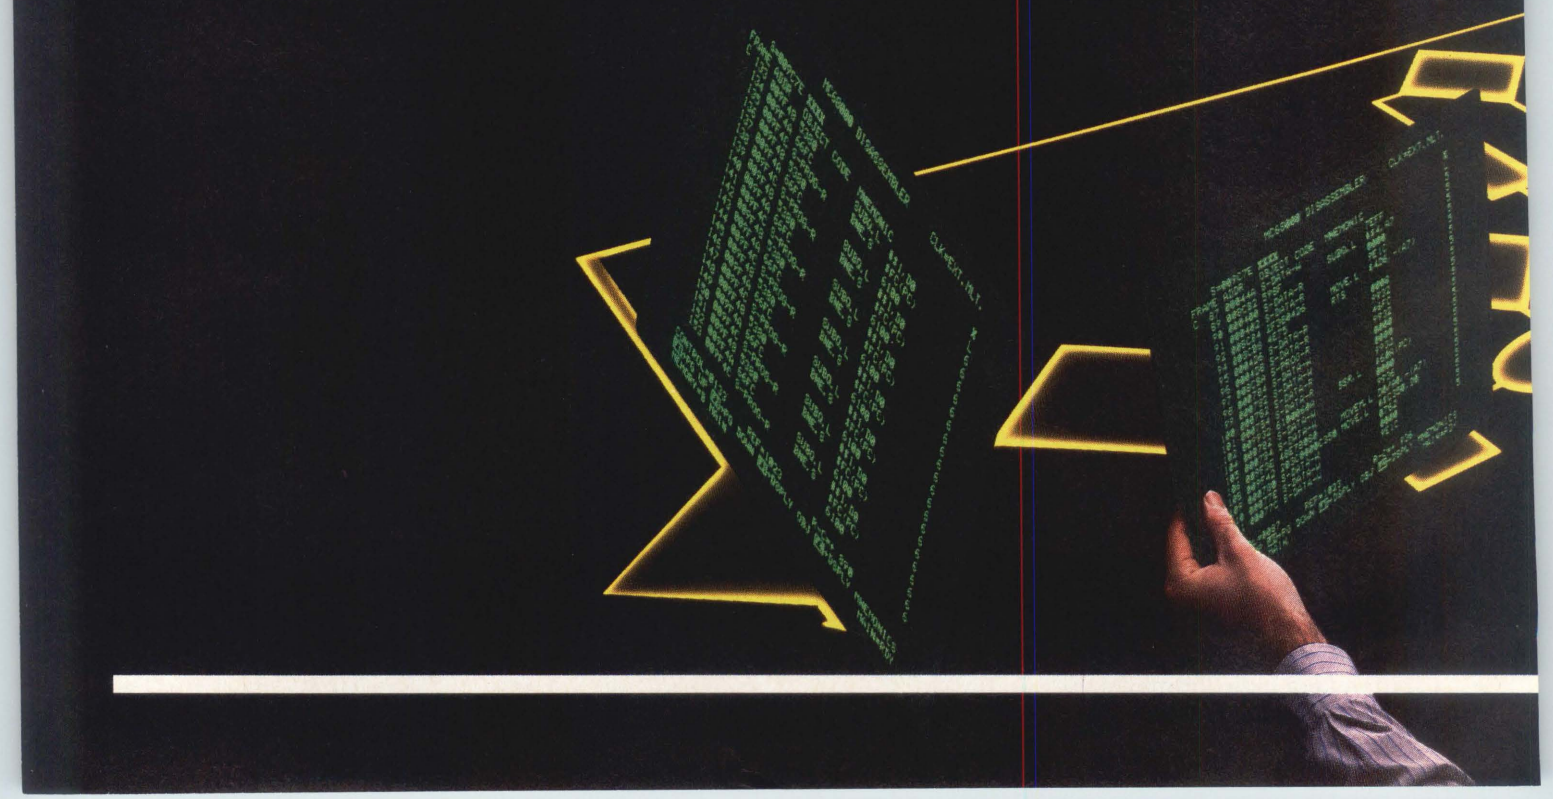
With the K105-D, you can define all your event patterns with symbolic labels. These labels can represent address locations, data values, fetched instructions or any control signal you're monitoring.

And since you create them, they're easier to understand and remember.

What's more, Trace Control uses English-like commands which make for shorter set-ups. Quicker comprehension. And a more productive engineer.

To make debugging even easier, add the optional dual floppy disk drive. With it, you can store up to 35 set-ups and associated recordings on a single diskette for instant recall.

Now, capture up to 4 segments of complex program flow to pinpoint problems fast.



You won't get caught in a loop.
By using the K105-D's loop counter, you can follow a sequence of events for a specific number of iterations. Then stop at will.

It's a quick way for you to test the repeatability of a program function. And verify that your code is having the predicted effect.

Integration made easy.
For the hard part (hardware/software integration), now there's a simpler solution. The K105-D's 100 MHz high speed link.

It's an extremely precise way to see how timing affects software execution. And vice-versa.

Again, for this task, Trace Control can help you track down hard-to-find faults. So you end up with a program that flows smoothly. Error free. In a fraction of the time it used to take.

Performance above all the rest.
At Gould, we engineer every instrument to be clearly the best in its class.

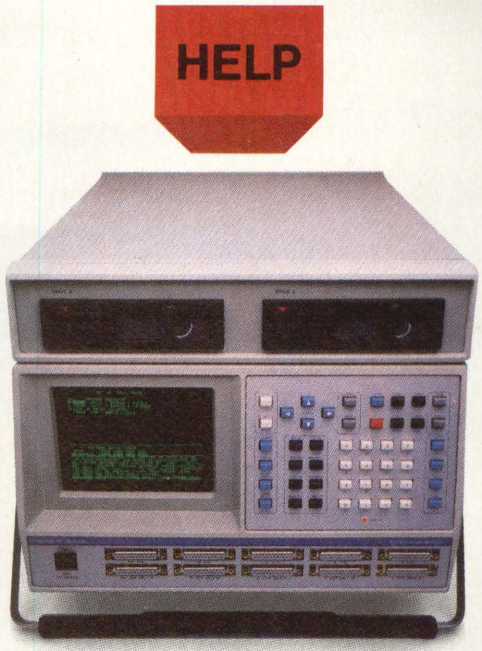
So now with your insight and the K105-D analyzer, you can capture the problem. And solve it. Fast.

For the name and number of your nearest Gould sales representative or for a detailed brochure and application notes, call us toll-free at **(800) 538-9320.** In California, call **(800) 662-9231** or **(408) 988-6800.** Or write Gould Inc., Design & Test Systems Division, 4600 Old Ironsides Drive, Santa Clara, CA 95054-1279.

United Kingdom: Gould Inc., Design & Test Systems Division, Hainault Operation, Roebuck Road, Hainault, Essex 1G6 3UE, United Kingdom. Phone: (44) (1) 500-1000.

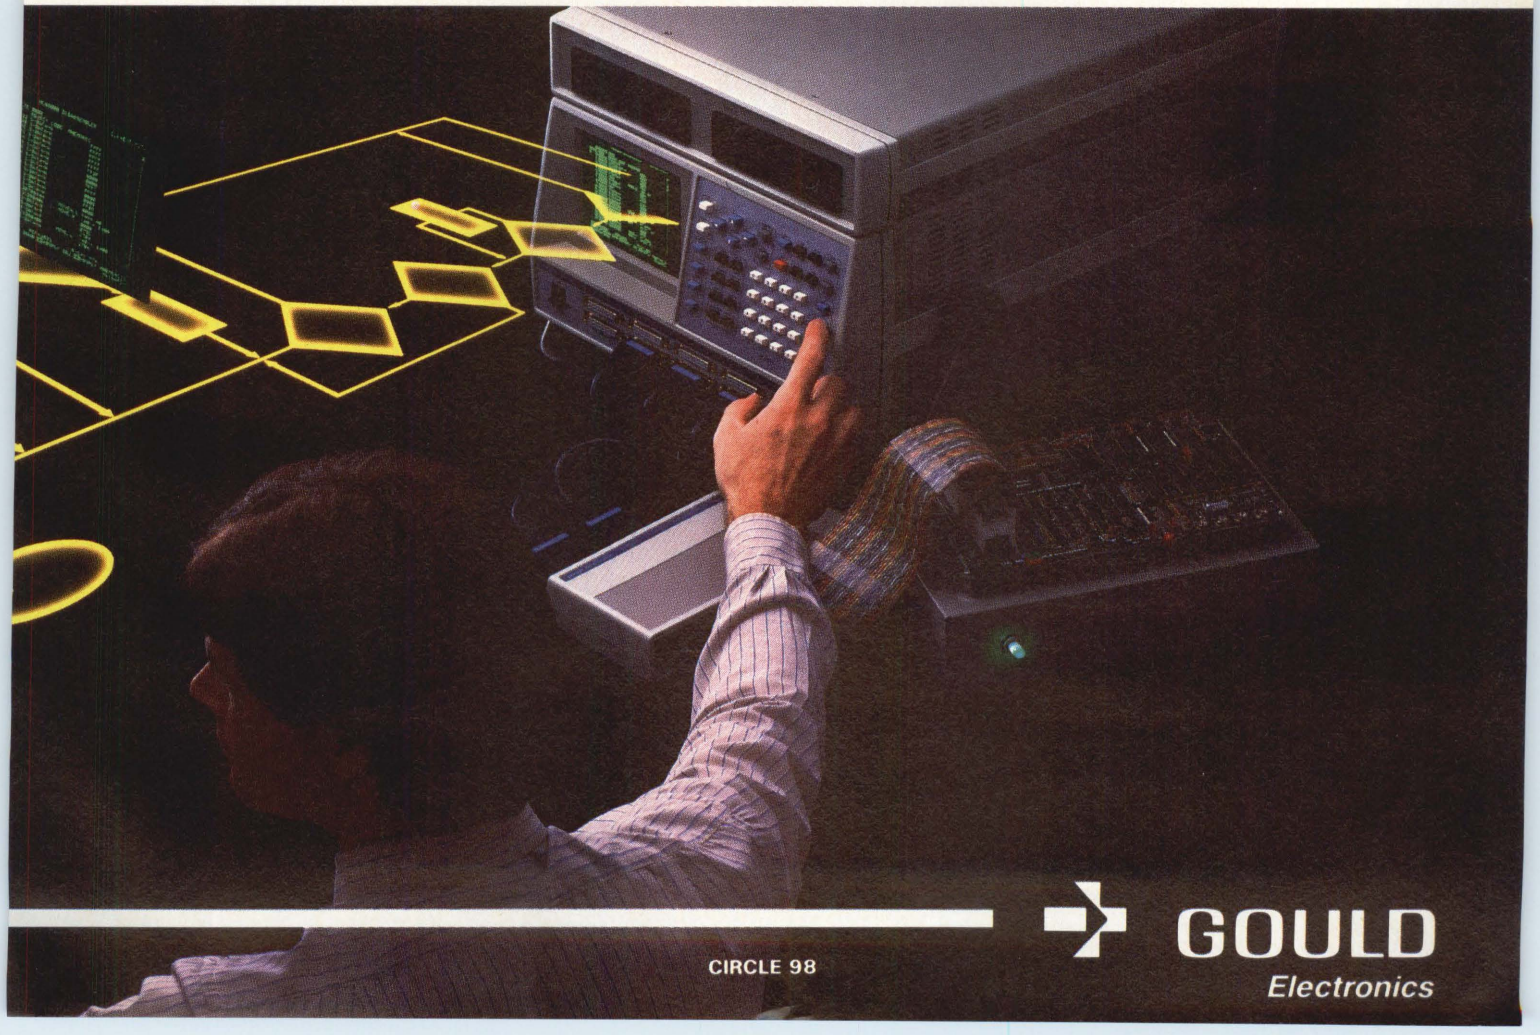
West Germany: Gould Inc., Design & Test Systems Division, Dieselstrasse 5-7, D-6453 Seligenstadt 3, West Germany. Phone: 06182/801-1.

™Trace Control is a trademark of Gould, Inc.
® Z80 is a registered trademark of Zilog, Inc.



The K105-D gives you two levels of HELP at the touch of a button. First, step-by-step operating instructions that appear along the bottom of the analyzer screen. Second, a menu that allows you to select more detailed "help" should you need it.

© 1984, Gould, Inc.



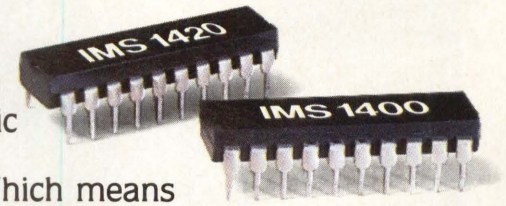
GOULD
Electronics

CIRCLE 98

In 16K static RAMs, INMOS choices...starting



means high-performance with 35ns access times



When you're looking for top memory performance in your system, look to today's leading supplier of high-speed static RAMs...INMOS. Because INMOS gives you a choice of speeds, organizations, power dissipation and packaging. Which means you can optimize your designs, without compromising on cost or performance.

High speed. The INMOS family of fast static RAMs deliver access times of 35, 45, and 55ns. Perfect for high-performance tasks such as buffer memories, graphics, and main-store memory. And they give you a low-cost alternative to expensive 35ns ECL parts. Choose our by-1 organization for deep memory. Or the by-4, which also offers versions with 30 and 40ns chip-select access times, for writeable control store and cache memory. Most versions are available in plastic DIPs, ceramic DIPs, and chip carriers.

Low power. For high performance combined with low power, INMOS offers plastic packaged SRAMs with 70 and 100ns chip-enable access times in by-1 and by-4 organizations. Both have maximum power dissipation of 495 mW active, 83 mW standby. The by-4 organization is also a better alternative to 2Kx8 SRAMs in many high-speed applications. Best of all, they're both low cost. Here's your answer for fast systems that are power critical and price conscious.

Military versions. When you need fast RAMs you can trust in hostile environments, INMOS can supply 55 and 70ns parts with specs guaranteed over the full military temperature range of -55°C to 125°C. These parts, available in side-brazed ceramic DIPs, chip carriers, and flat-paks, conform to MIL-STD 883B processing, including screening to class B Method 5004 and quality conformance to Method 5005.

Check the chart. With all these choices, its easy to see why INMOS is the leading high-speed RAM supplier. Look for the part you need, then call an INMOS distributor for complete details.

Organization	Speed (ns)	Power (mW) Active	Standby	Part No.
16Kx1	35	660	110	IMS1400-35
	45	660	110	IMS1400-45
	55	660	110	IMS1400-55
	70	495	83	IMS1400-70L
	100	495	83	IMS1400-10L
4Kx4	45	605	165	IMS1420-45
	55	605	165	IMS1420-55
	70	495	83	IMS1420-70L
	100	495	83	IMS1420-10L

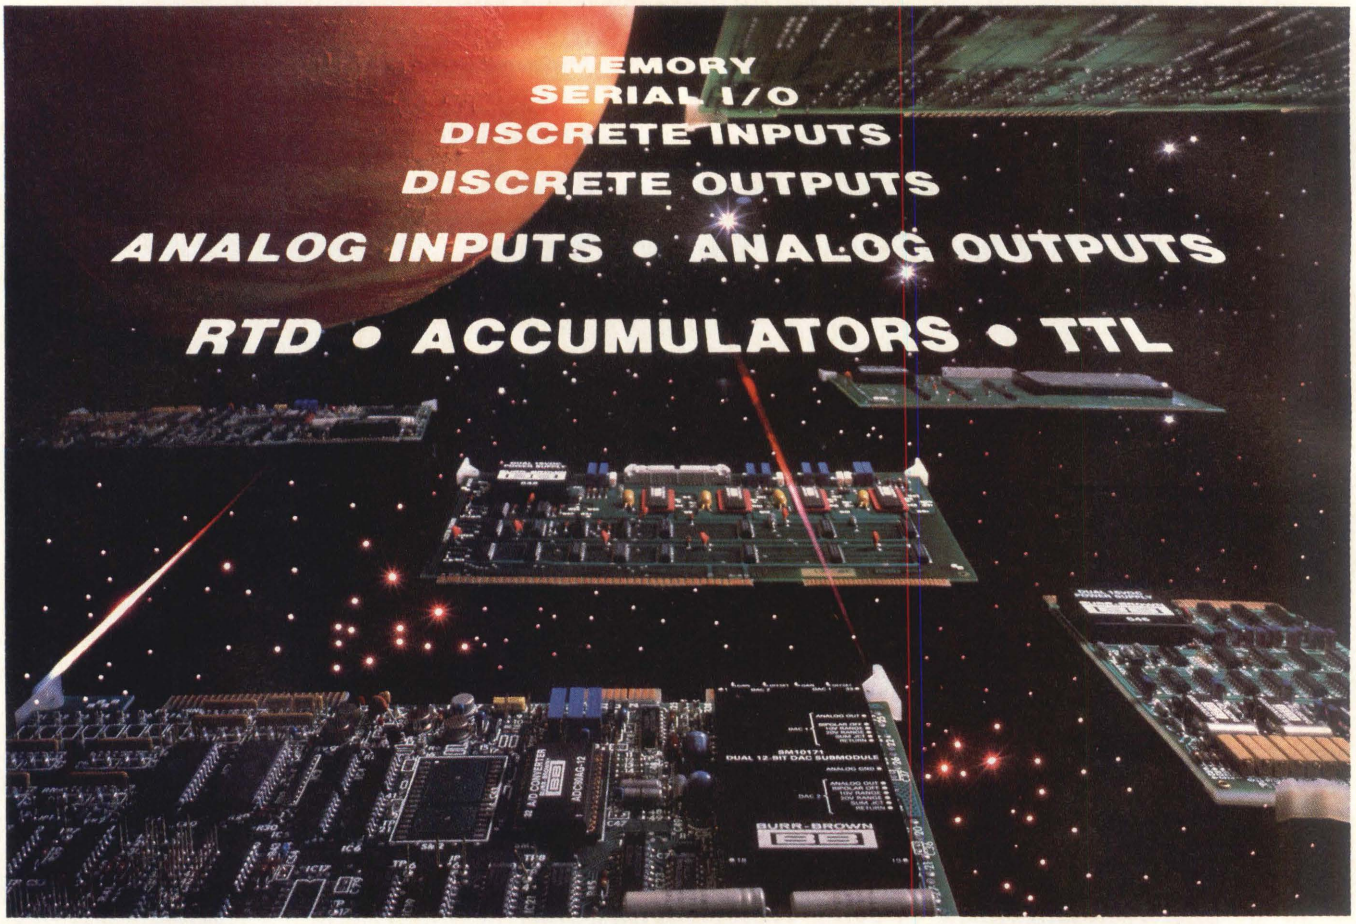
INMOS Distributors: Anthem Electronics, Arrow Electronics, Falcon Electronics, Future Electronics, Lionex Corp.

inmos

P.O. Box 16000 • Colorado Springs, CO 80935 (303) 630-4000 TWX 910/920-4904 • Burlington, MA (617) 273-5150 • San Jose, CA (408) 298-1786 • Torrance, CA (213) 530-7764 • Minneapolis, MN (612) 831-5626 • Baltimore, MD (301) 995-0813 • Dallas, TX (214) 669-9001 • Atlanta, GA (404) 475-0709 • Whitefriars • Lewins Mead • Bristol BS1 2NP • England Phone Bristol 0272 290 861 • TLX: 444723

inmos and IMS are trademarks of INMOS' Group of Companies.

CIRCLE 99



A Universe of I/O Solutions for Multibus™ Systems

Getting signals to and from your Multibus™ is our business. We've already done the design and engineering work to solve many of your microcomputer interfacing problems, so you can spend more time on total system integration and performance. We have over 40 Multibus™ analog, discrete, and serial I/O cards "on the shelf," ready to go to market in your systems right now.

These low cost boards give you the flexibility you need for a wide range of applications.

Benefits include:

- fast, *simple* programming
- memory or I/O mapped
- easy installation
- powered from system bus
- flexible system configuration
- fast system development

Signal Types

- ±10V
- ±10mV
- 4 to 20mA transducers
- thermocouples
- RTD
- DC sense
- AC sense
- contact closures
- TTL
- discrete I/O
- pulse counter
- pulse width
- frequency
- period
- RS-232-C

Write or call for complete specifications and pricing information.

Data Acquisition and Control Systems Division
3631 E. 44th Street, Tucson, AZ 85713 • (602) 747-0711

BURR-BROWN®

Putting Technology To Work For You

ALABAMA, (205) 663-2831 - **COLORADO**, (303) 771-1530 - **DELAWARE**, (302) 478-8211 - **FLORIDA**, Tampa, (813) 885-1796; Jacksonville, (904) 268-1468 - **GEORGIA**, (404) 449-3382 **ILLINOIS**, (312) 579-9300 - **INDIANA**, (317) 257-7231 - **LOUISIANA**, (504) 887-8550 - **MASSACHUSETTS**, (617) 339-5522 - **MINNESOTA**, (612) 835-2044 -**MISSOURI**, (314) 647-8100 **NEW JERSEY**, (201) 785-4450 - **NEW YORK**, Albany, (518) 869-3355; Buffalo, (716) 874-4900; Rochester, (716) 323-1250; Syracuse, (315) 455-2469 - **OHIO**, Northern, (216) 247-3295; Southern, (513) 772-5544 **OREGON**, Medford, (503) 772-3187; Portland, (503) 283-2555 - **PENNSYLVANIA**, (412) 366-6664 - **TENNESSEE**, (615) 689-8513 - **TEXAS**, (713) 491-2000; Beaumont Area, (713) 842-5950 - **UTAH**, (801) 566-9251 **WASHINGTON**, Richland, (509) 943-6664; Seattle, (206) 762-6050 **WISCONSIN**, (414) 352-5777 - **CANADA**: (403) 458-4669; (416) 625-0600; (514) 774-5829; (604) 434-2611; (613) 725-0288

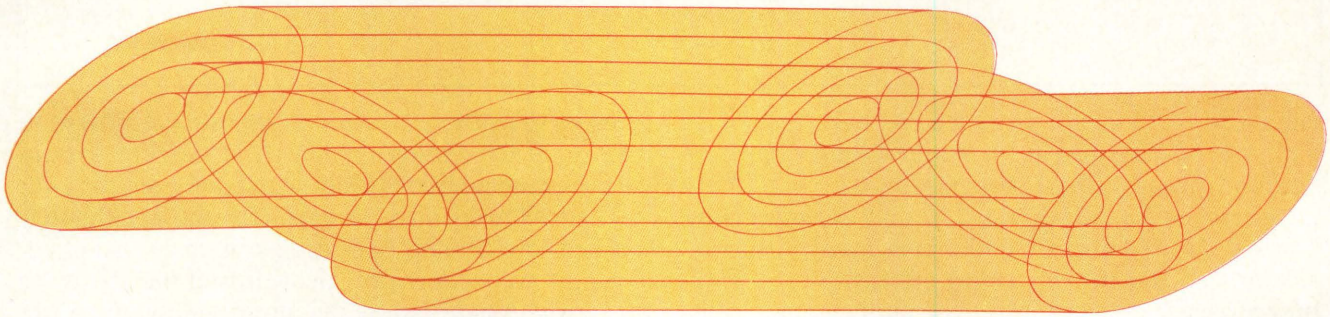
Other areas: **Burr-Brown,** 3631 E. 44th Street, TUCSON, ARIZONA 85713 USA - Tel: (602) 747-0711 - TWX: 910-952-1115 - Telex: 66-6491 - Cable: BBRCORP

Multibus™—Intel Corp.

CIRCLE 137

DUAL PORT SOLVES COMPATIBILITY PROBLEM

A single operating system incorporates both System V and 4.2 BSD Unix versions to solve program development/ portability problems.



by Ross A. Bott

Historically, Unix has attracted followers because of its power, flexibility, and portability. Because of AT&T's initial reluctance to market Unix directly, two Unix versions have become popular for superminicomputers. The first, 4.2 BSD, is the latest of the popular Unix enhancements from the University of California at Berkeley. The second is AT&T's own officially supported Unix System V.

With AT&T's renewed interest and strong support, System V and its successors are attractive as the probable long-term standard, even though the presently written System V lacks some of 4.2 BSD's desirable features, such as virtual memory support and enhanced file access. Unfortunately, while the two systems are similar in many respects, they are not user and application code compatible.

Ross A. Bott is a project manager at Pyramid Technology Corp, 1295 Charleston Rd, Mountain View, CA 94043, where he is responsible for networking R&D planning. Mr Bott holds a BS in mathematics and psychology from Stanford University and a PhD in cognitive psychology from the University of California, San Diego.

Until now, computer system integrators, designers, and end users had to choose between the two operating systems. Now, however, the dual-port OSx operating system from Pyramid Technology Corp offers separate Unix "universes" in which users can move easily from one environment to the other.

Too many Unix versions

Because the Unix community has not had a single controlling force, there are almost as many nonstandard versions of Unix as there are types of computers running it. Trying to meet perceived needs, system builders select one of the versions and enhance it where necessary to provide what seems to be a reasonable subset of advanced features. Some system builders have adopted one Unix version, modified it slightly, and then staked a claim to generic portability. Yet, this portability is limited to other versions of the same Unix flavor only. Likewise, vendors of proprietary operating systems have sought to appeal to Unix enthusiasts with versions that are grafted on top of an existing, proprietary operating system. In this way, they claim portability while maintaining their own operating system.

As might be expected, this structure leads to inefficient Unix execution and a decrease in performance. Furthermore, portability depends on which Unix version is supported and how many standard or

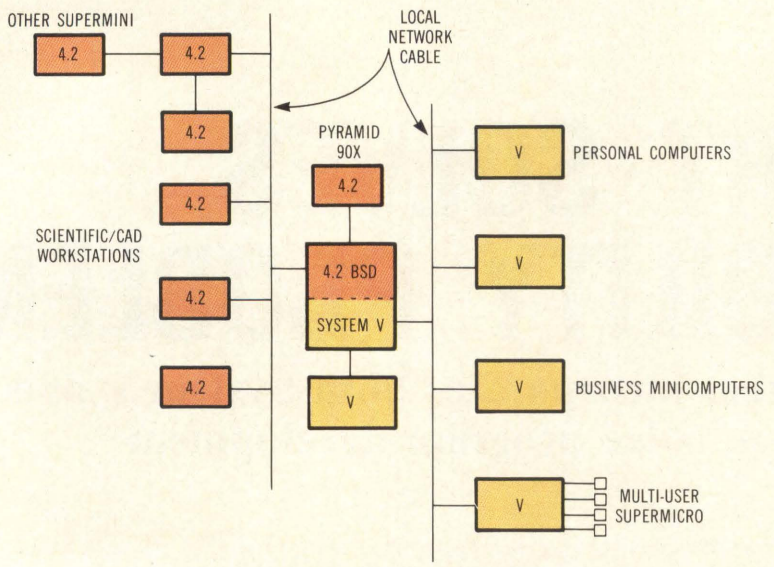


Fig 1 As networked, distributed operating environments become more common, user-transparent linking of incompatible System V and 4.2 BSD Unix operating systems will be typical. OSx-based systems such as the Pyramid 90x superminicomputer facilitate this linking.

enhanced features are provided. At best, true portability is gained only at the expense of overall system and program performance.

Others have tried a third route, selecting one version of Unix, and then adding as many attractive features of the other version as possible. Performance varies depending on the skill of the implementers, and portability is constrained by the patchwork of added features.

Forcing system integrators to choose between popular Unix versions or to make their own by patching the two together is unfavorable. The 4.2 BSD version offers virtual memory support, superior disk I/O, and the networking support necessary to run Unix effectively on large machines. On the other hand, System V provides a variety of time-tested tools to run Unix in commercial environments. With AT&T's backing, it should provide better support and possibly better potential for future development.

While it is likely that the two systems will eventually evolve into a single standard, it is also likely that the double standard will exist for several more years. In addition, it is impossible to say just which operating system will emerge victorious (if any), and what form it will take. The OSx attempts to solve the problem by consolidating both leading standards into one comprehensive system that not only solves the compatibility problem between the two, but also provides the means to support future versions in an equally compatible way.

Defining system goals

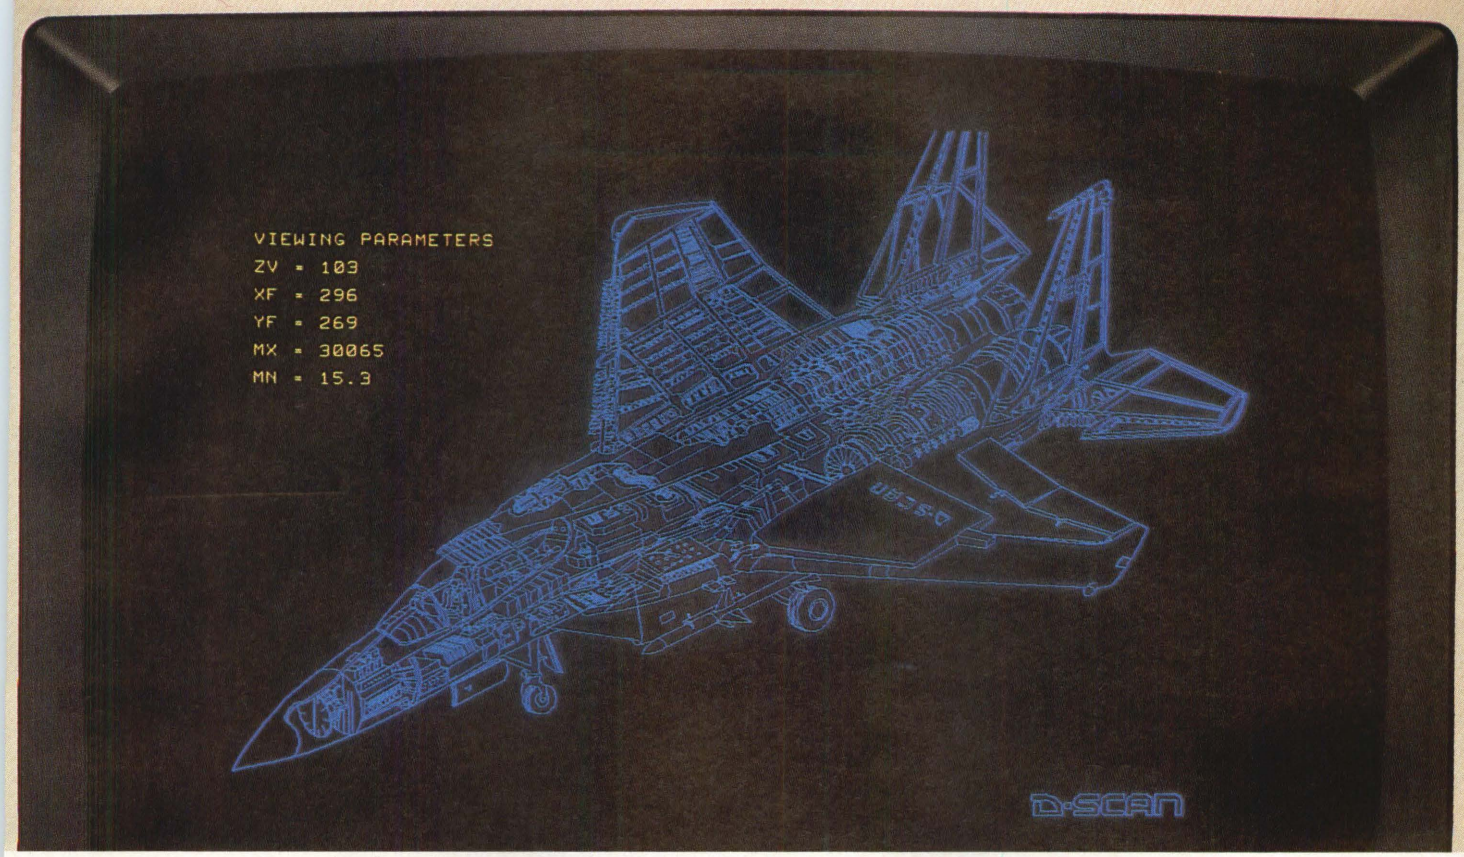
The OSx aims to provide clean, standard versions of both 4.2 BSD and System V, while enabling users in one environment to take advantage of special features in the other. It does so by allowing users to choose, at log-in time, either the 4.2 BSD or System V environment and to easily switch between the two.

The system is also designed to support the networked, distributed operating system environment expected to become prevalent during the 1980s. In this environment, some machines will want to run System V for commercial users, while others will want 4.2 BSD for technical or scientific applications. The OSx environment provides a gateway enabling a network to be established between these incompatible operating systems (Fig 1). Allowing programmers (who usually prefer the enhancements of the 4.2 BSD version) to develop their programs in that system's technical environment, even if their target systems will be System V-based commercial systems, is another OSx objective.

Finally, the operating system has a permanent set of hooks or pointers to support future versions of either operating system. Thus, it can remain compatible with the latest revisions from either source.

In contrast to the additional software processing overhead, which must be suffered when a layering technique is used to place one operating system software interface on top of a primary interface, the dual-universe approach has no performance penalty. In a layering technique, input data or commands are reformatted into acceptable form for the underlying calls or commands. This data conversion reduces system performance by as much as 40 to 50 percent.

The hybrid technique, wherein additional features are added directly to a Unix kernel, extracts little performance penalty. But, the compatibility issue remains unsolved. In reality, a hodgepodge of extra



High Speed Craft.



No, not the fighter. The craft we're talking about is yours. Because with the D-SCAN GR-2414, you'll experience something you've probably never felt before. High speed response.

Gone are the days when you had to sit around waiting for the host computer to answer your commands.

Quickly, here's the story. Speed. Efficiency. Power. Extraordinary visual acuity. And local control!

We're not done yet, either. You can display up to 1024 colors (not the usual 256). Read them flawlessly on our 1280 x 1024 display (with 60Hz non-interlaced refresh rate!). And expand the 4 bit planes to 10.

The 2414 also gives you display list storage that's simply unmatched in its price class. Starting with 256K and field-expandable to over 1Mbyte! That's two, three, sometimes even four times more storage than you can find in so-called competitive terminals.

Then, to make sure you get the picture you want, the 2414 gives you local picture manipulation. Local transformations. And local graphics tablet support. Which means you'll have true local interactive support. You won't get bogged down while the host spews out accounting reports. And you get instant visual response every time you need it.

There's more to those pictures, too. More complexity. More detail. More information on the 20" screen at any one time. Plus, the handy ability to manipulate the image in more ways. And display alphanumerics with beautiful precision.

A very pretty picture indeed. What's more, you get easy interface to our color hardcopier. Programmable function keys. A user-definable look-up table. And user-definable hardware anti-aliasing.

But hardware technology notwithstanding, you also get a powerful software story with the 2414. For instance, it's compatible with the TEK 401X instruction set and Plot 10 package. So you can save a lot of time and aggravation implementing graphics software. Upgrade quite painlessly. And eliminate costly re-learning curves.

There's something heavyweight to remember, too. Our service. We design, build, sell and service all of our equipment. So you never have to wander around wondering who to call for help. And waits are lifted because you get direct service from 13 offices across the U.S.

So call today to learn more about the GR-2414. We'll get back to you with utmost speed. Just call your local Seiko Instruments sales person, or us at (408) 943-9100. Write 1623 Buckeye Drive, Milpitas, CA 95035.

You'll find there's a new high speed craft that's drawing a lot of attention. Yours.

SEIKO INSTRUMENTS

D-SCAN is a trademark of Seiko Instruments & Electronics, Ltd.
TEK and Plot 10 are registered trademarks of Tektronix, Inc.
© 1984 Seiko Instruments U.S.A., Inc., Graphics Devices & Systems Division.

CIRCLE 101

features gains nothing but another unwanted, nonstandard Unix version.

OSx provides a high performance operating system interface and maintains strict compatibility for both users and programs. Since many users desire only the features of their preferred Unix and do not necessarily want their operating system to coexist with another version, there is a clean interface to both systems with no performance penalties.

Regarding implementation

In this solution, the key idea is to have separate Unix universes. Both 4.2 BSD and System V share many features, but have distinct twists in some system calls and user command responses. The OSx dual universes are bridged through a defined and accessible, though invisible, universe flag in its kernel. Users need only be aware of the particular Unix universe that they prefer: commands, libraries, header files, and system calls appear exactly as expected.

But, as shown in Fig 2, in each directory for which 4.2 BSD and System V have different contents (for example /bin), two OSx directories exist. Thus, / .ucbbin and / .attbin both exist in OSx and correspond to the 4.2 BSD and System V universes, respectively.

When the user enters a desired directory or command (in normal input syntax), the kernel checks OSx's invisible universe flag to determine which universe is in use, then finds the appropriate directory or command. Since the directory or command matches what is expected from 4.2 BSD or System V, compatibility is ensured and the user sees the expected response from any command.

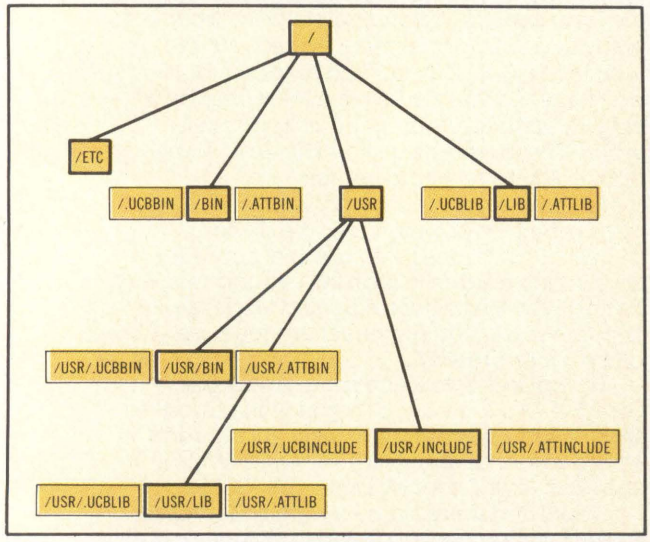


Fig 2 The dual-universe directory structure underlies the normal transparent directory structure perceived by the user and/or the programs. Bold lines indicate what the user sees; lighter boxes reveal underlying directories with universe-indicating attribute fields.

The time required to make the universe check is minimal and operating system performance is not impaired. And, since many of the functions and calls of both operating systems are identical, redundancy has been avoided by providing only one OSx version.

Since the 4.2 BSD version is faster in many respects than System V (because of its faster file access, demand-paged virtual memory, and a generalized "socket" mechanism), OSx utilizes 4.2 BSD internally. However, if future versions of AT&T's Unix offer better internal features and performance, they can be incorporated into OSx.

As mentioned earlier, the two separate 4.2 BSD and System V universes coexist in the same file structure and share a common kernel. At any given time, users operate in only one universe but can easily switch to the other. The initial universe is determined at log-in time by an entry in a /etc/u_ universe file maintained for each user. At any time, the user can enter the alternate universe by typing one of two commands:

warp att (or just att) — to enter System V
warp ucb (or just ucb) — to enter 4.2 BSD

These commands fork a shell within the alternate universe. Using them is similar to typing csh while using sh in conventional Unix. Thus, the operating system user can return to the original universe by typing ^D. In addition, typing the word "universe" will list the universe in use. OSx also allows users to set and switch the universe in which they are operating, using the ucb or att commands.

Users can easily access a command in one universe while operating in the other by prefixing the command with the appropriate prefix, either att or ucb. Upon finishing the command, the system returns to the original universe. What is more, inputs and outputs of different universe commands can be piped to each other with wild card expansion and redirection if desired. Thus, users in one system gain the best of both operating systems while losing none of the benefits of either.

Multiple levels

Although the Unix operating system is not as modular as it might be, the kernel can be divided into a set of concentric levels, as shown in Fig 3. The innermost level concerns itself with virtual memory and process management, I/O device control, and interfaces with the architecture of the host machine (see Panel, "Inside the system architecture"). Ninety-five percent of the changes to port a Unix kernel to a given architecture are isolated to this level.

On the other hand, the particular version of Unix interface supported, such as 4.2 BSD or System V, is

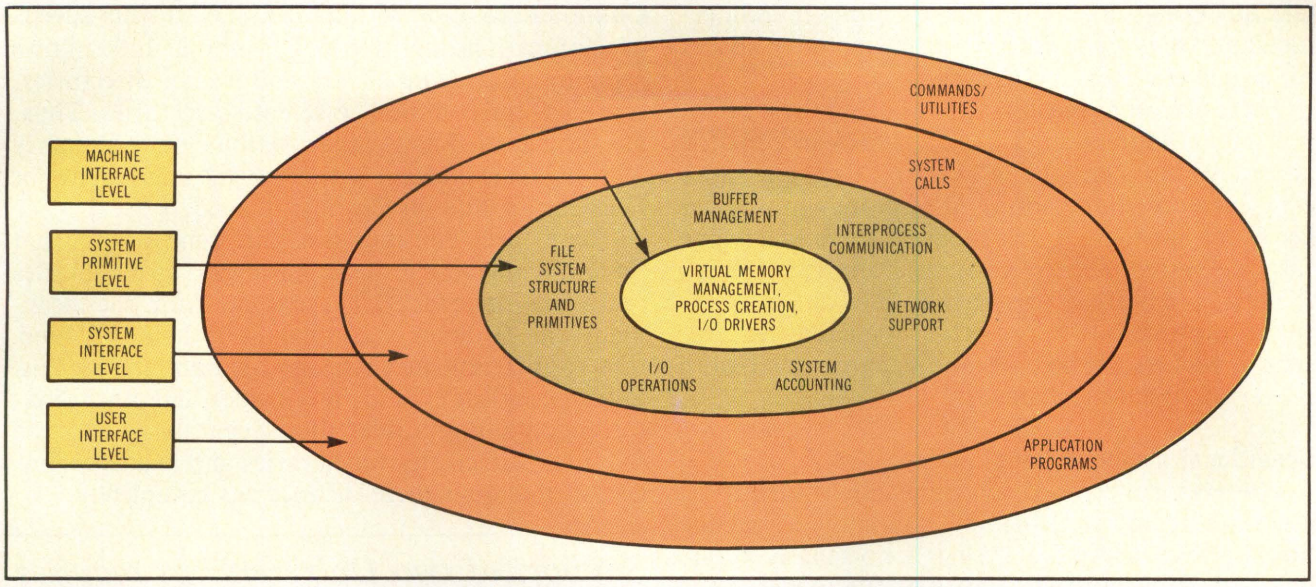


Fig 3 Unix is divided into levels that separately handle user, system, services, and machine interface functions. Calling conventions and software structures isolate levels from one another and allow modular software changes to meet new needs.

in essence defined by the system interface level, which directly handles system calls. In between these two levels is the system services level, which defines and provides the large scale software facilities to support the system calls. The structure of these facilities is, to some extent, what distinguishes Unix from other operating systems.

From this perspective, it is clear that simultaneous support of two Unix versions requires major changes in the system interface level. Some modifications may be required of the system service level; these changes tend to be the addition of new features, rather than drastic overhauls. For example, System V's semaphores and messages impact this layer. Virtually no changes must be made to the machine interface level.

Since the Unix version chosen rarely impacts the inner levels, a fairly strict 4.2 BSD implementation with the associated performance gains for the inner levels (Fig 4) has been selected. This can be done and both 4.2 BSD and System V can still be supported, if the System V interfaces in the outer level are fully supported.

The 4.2 BSD version was selected for internal operations for several reasons. First, it is the only Unix version to support the virtual memory and demand-paging algorithms found on larger machines like the 90x. In OSx, 4.2 BSD has been enhanced to

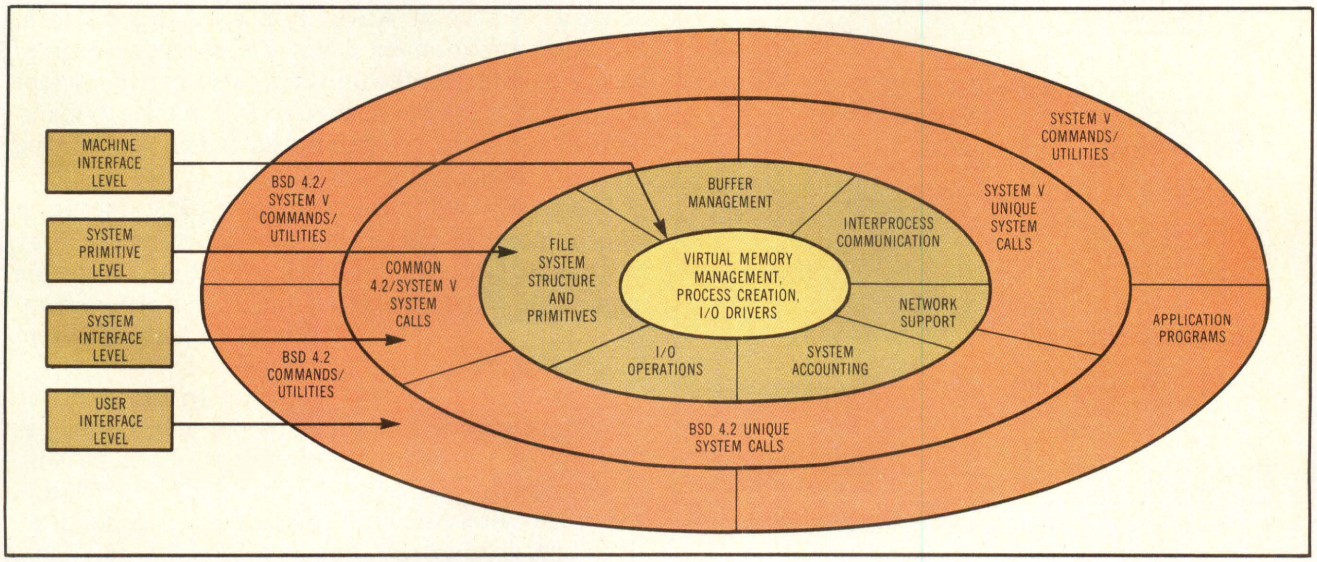


Fig 4 In OSx, the differences between System V and 4.2 BSD are handled in the outer three levels. Many user interface conventions are the same; some unique commands and utilities have been added. Both operating systems' call handling functions exist at the system call level.

take advantage of a reduced instruction set architecture and to provide larger virtual memory spaces for processes—up to 4 Gbytes.

There are other reasons for choosing the 4.2 BSD. For one, it addresses the I/O bottleneck in Unix file access. With its larger file blocking sizes and more efficient access algorithms, it is more suitable for use with a large, multi-user superminicomputer where fast file system access is a must, and its network facilities are more extensive than that of System V. In addition, it has many other attractive features such as sockets, flexible length file names, and more powerful signal mechanisms.

Conditional symbolic links

The OSx concept of alternate environments is accomplished by implementing conditional symbolic links. These work like 4.2 BSD's normal symbolic links, except that the directory or file to which the symbolic name points is dependent upon which universe is in use. For example, to set up the /bin directory when the two universes are ucb and att, two separate directories are created (/.ucbbin and /.attbin). Then a system developer uses the following command:

```
In -c ucb=/ .ucbbin att=/ .attbin /bin.
```

When using this procedure in the 4.2 BSD universe, /bin will look identical to / .ucbbin, both for executing commands and listing the directory. These underlying directories will typically be invisible to users listing directories, if the name is preceded by a period.

Conditional symbolic links work as fast as symbolic links and hardly influence performance.

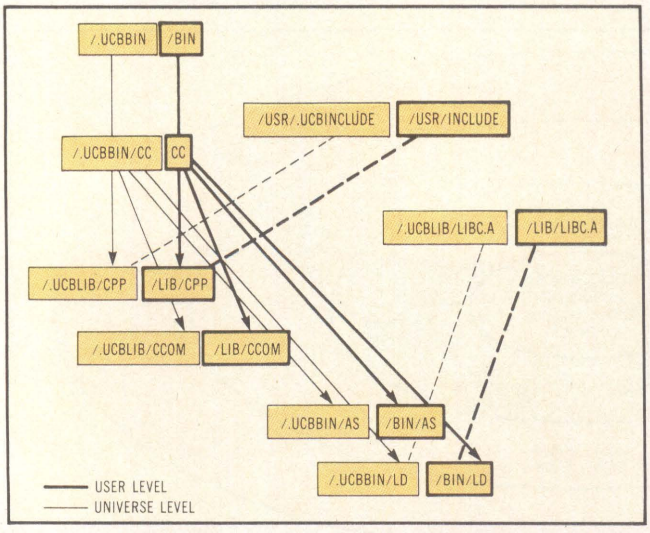


Fig 5 Because programmers can access both universes, it is important that application programs be fixed to one environment for compatibility. A program execution environment is established by linking compiler, libraries, and header files during compilation/linking.

Command hashing within the C shell further speeds execution. The conditional symbolic link allows separate underlying library and header directories, as well as the common command directories. Thus, as shown in Fig 2, the /usr/include directory will in fact address either /usr/.ucbinclude or /usr/.attinclude directories of headers.

There are considerable differences in what System V and 4.2 BSD programs expect from libraries and kernel interfaces. These range from the meaning of the different return values in the printf function to different mechanisms for handling signals. Similar conflicts exist in /usr/include and other directories containing header files. Separate 4.2 BSD and System V directories for libraries and headers allow OSx to maintain these distinctions completely.

Support of two Unix versions requires changes in the system interface level.

The correct header files and libraries are automatically accessed by invoking a distinct compiler such as cc or f77 within the desired universe. The various portions of the compiler (such as cpp or ccom) live within the same universe as the compiler itself. The header and library directories are also included/linked from that universe (Fig 5). Once the program has been compiled, its behavior with respect to 4.2 BSD or System V compatibility has been established and it can be executed in either universe.

Program execution environment

The program execution environment is a characteristic of each program running on top of the OSx kernel and can be contrasted with the universe concept in several ways. For one, the program execution environment is not file system or process based. In contrast, the universe characteristic is kept on a per process basis, and is used during file system access only to determine which underlying directory structure is relevant.

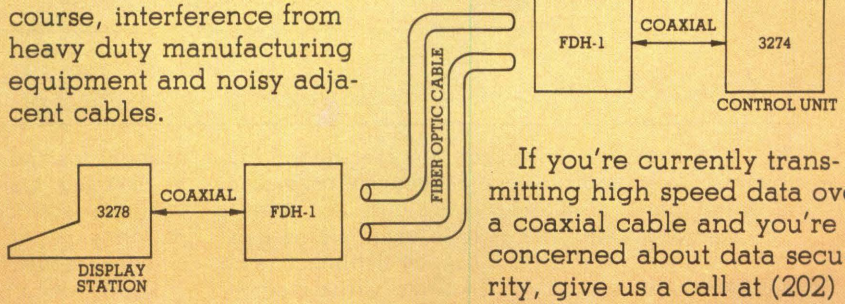
There are two more contrasts. The execution environment of an object file is determined by the header files and libraries used to compile the application. And, the kernel utilizes this program environment to provide differential 4.2 BSD or System V behavior where they conflict.

A good example of the distinction between universe and execution environments may be seen in application software development. It is possible to develop software for a particular Unix version, using the alternate universe. Thus, if a software developer prefers the tools provided by the 4.2 BSD environment, but is developing an application package to be transparently portable to a System V environment, all of the designing and editing can

Protect your sensitive IBM data with our new Fiber Optic Link

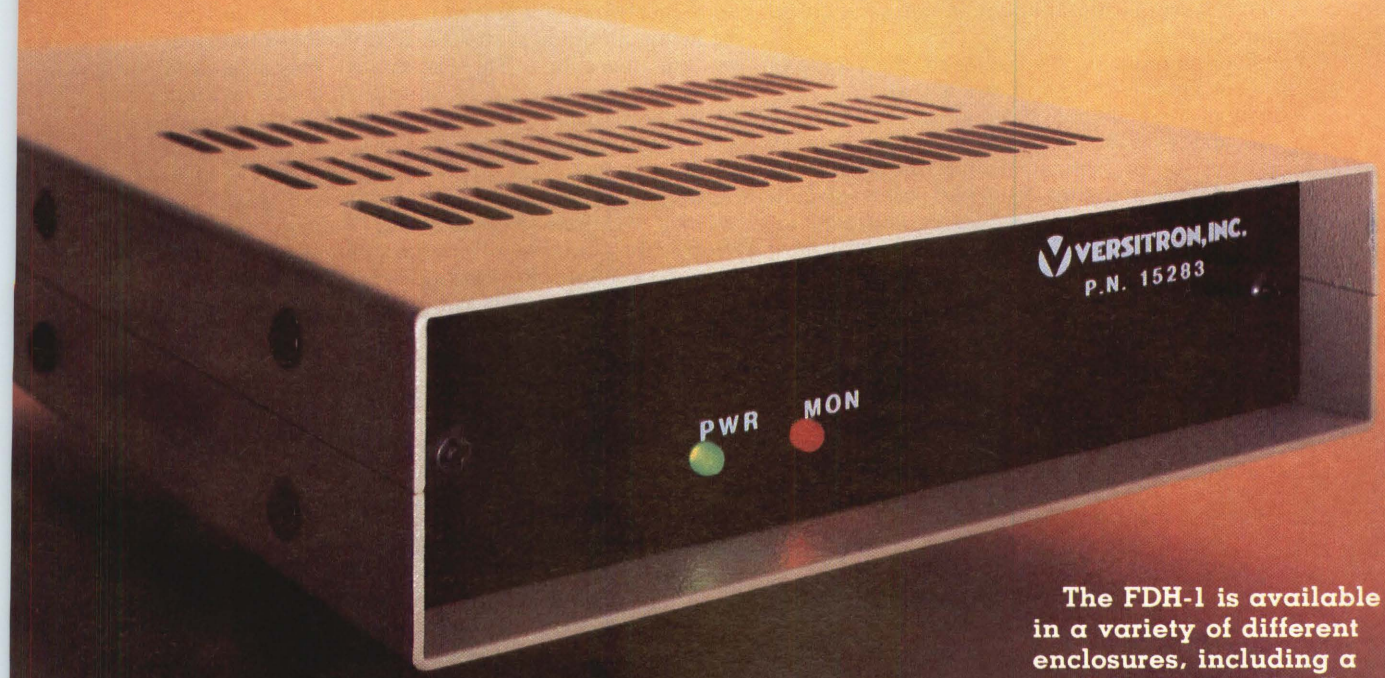
- **Plug compatible with IBM series 3250, 3270A, and 3270B equipment.**
- **Replaces coaxial cable with fiber optic cable.**
- **Up to 1 Km operating range — virtually immune to electromagnetic interference.**

Versitron's FDH-1 (fiber optic digital hybrid) was designed to replace the coaxial transmission path in systems equipped with the IBM 3250 or 3270 series equipment. The simple installation of a fiber optic link provides two very important benefits to the user. First of all, the security level of the transmission link is greatly improved since fiber optic cables are inherently immune to conventional wiretapping techniques. Secondly, the system operating capability will be enhanced since fiber optic cables are impervious to virtually all types of electromagnetic interference. This includes, of course, interference from heavy duty manufacturing equipment and noisy adjacent cables.



Versitron's FDH-1 combines the high speed capabilities of a coaxial cable with the inherent advantages of a fiber optic cable. By interfacing directly to the coaxial cable, the FDH-1 appears totally transparent to the rest of the system; thus eliminating any operating restrictions.

If you're currently transmitting high speed data over a coaxial cable and you're concerned about data security, give us a call at (202) 882-8464 and get all of the details on how our FDH-1 will not only protect your data; but may also actually increase the operating efficiency of your entire system.



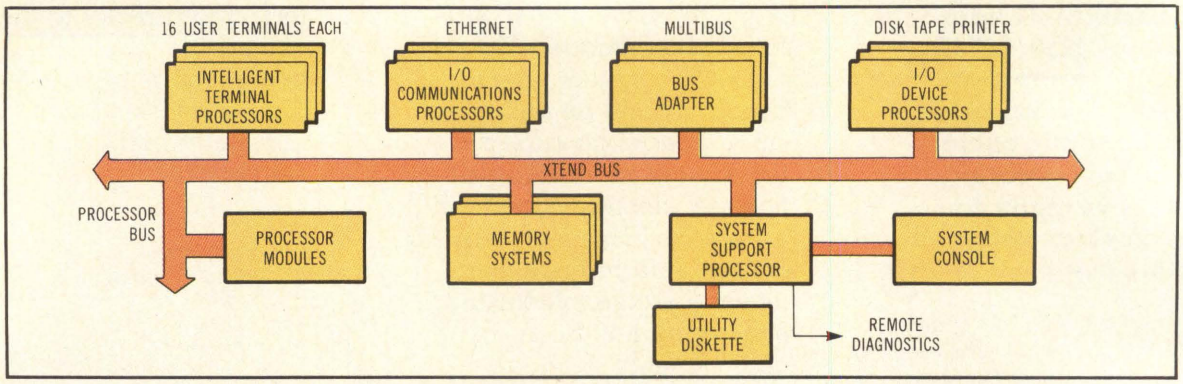
The FDH-1 is available in a variety of different enclosures, including a sealed unit specifically designed for EMI/RFI suppressed applications.

VERSITRON, INC.

6310 Chillum Place, N.W., Washington, D.C. 20011, TEL: (202) 882-8464
TWX: 710-822-1179

CIRCLE 102

Inside the system architecture



The OSx system is designed to execute on the Pyramid 90x superminicomputer system. This system is based on a 32-bit, Schottky-logic CPU with a 125-ns cycle time. The system incorporates local cache memory, pipelined overlapped instruction fetch and execute functions, and a 4-Gbyte virtual memory address space to support up to 128 users (see Figure). Intelligent I/O processors optimized for the I/O-intensive Unix environment speed data transfers at an average 220 ns per word.

Processor, memory, I/O processors, and system support processors communicate over a proprietary message-based bus. This xtend bus has a 32-bit data path and a 32-Mbyte/s bandwidth. On the bus, I/O requests are standardized through the use of intelligent I/O processors (IPOs) that act on message-based I/O commands.

In addition to reducing CPU overhead, the exclusive use of message-based I/O and IOPs allows future expansion by giving additional symmetric or asymmetric CPUs the ability to share memory and peripherals in a tightly coupled fashion. Finally, the xtend bus accepts industry standard Multibus (IEEE 796) controllers, peripherals, and accessories through a microprocessor-controlled Multibus I/O channel adapter.

The 90x system is optimized for efficient execution of high level languages through the use of Reduced Instruction Set Computer (RISC) techniques and a register-intensive architecture. The 90x is based on research performed at the University of California at Berkeley, Carnegie-Mellon, Stanford, and other computer science institutions that showed that the performance of RISC-type computer designs can surpass traditional computer architectures.

The 90x incorporates an instruction set that optimally places operands in registers so that memory references are minimized. To do this, the 90x uses 528 registers, each 32 bits long, implemented in stack form with 16 levels of 32 registers each. It also has 16 global registers. A register window makes a subset of the registers at each stack level accessible to the next level.

As a result, parameter passing does not require any data movement. The register stack also provides a uniform mechanism for procedure calls; 16 call levels can be accommodated without memory access for save-restore operations. The Fortran-77, Pascal, and C compilers can take full advantage of these features by generating code that keeps most of the procedure parameters, local variables, and temporaries within the register stack.

Networking is supported via standard Unix uucp utilities or Ethernet connections. For Ethernet networks, the company offers the Internet TCP/IP-based Networking Software Package (NSP) and the Ethernet Link Controller (ELC). Together, they allow multiple interactive links between 90xs and a variety of Unix-based computers and workstations.

be done in a ucb universe. The program is then compiled using

```
att cc -g program.c -o program.
```

After compilation, the program can be debugged using either

```
att sdb program
or
dbx program.
```

Note that in the latter case, dbx runs within the ucb universe, while the program itself runs within the att universe. This is possible because the system calls and header file assumptions are all made at compilation time. The kernel can then recognize that it needs to handle the System V or 4.2 BSD system calls when they occur.

In some cases, it is possible for OSx to mix 4.2 BSD and System V code within libraries or headers without affecting either system interface. However, whenever there is even a remote possibility that combining the code will produce a possible incompatibility within one of the universes, separate versions should be maintained.

Where portability is not affected, OSx supports only one format. For example, OSx supports a single object format (BSD style) and supports full flexible-length symbols in either environment. This allows debuggers from either universe to be used on an object file from either universe. In order for application developers to guarantee portability, there is a flag that can be passed to the C compiler within the System V world. This flag will warn the user

about systems that are not unique to the first eight characters—a characteristic of System V symbol names.

In the dual port of both 4.2 BSD and System V kernels, several software design issues concerning combining, merging, or duplicating features require considerable thought and understanding of the eventual effect on end users and programmers. Among the easiest changes to implement are system call translations, separate system call entry points, superset structures in system header files, and the addition of new, independent modules like System V semaphores and shared memory.

It is sometimes possible to translate incompatibilities (manifested through slight conflicts in data structures or values returned from system calls) from one operating system version to the other. Thus, the layered approach should only be taken when no significant loss in efficiency is involved. Such layering can occur at either the system call stub within libc (the C library) or within the system interface level. Because of the 4.2 BSD basis of the OSx kernel, all of the translations are from System V to 4.2 BSD. Some examples include the calls time, ulimit, and dup.

In terms of signal types, 4.2 BSD is nearly a superset of System V, and offers the starting point of OSx code.

Separate 4.2 BSD and System V system call entry points are a convenient method of providing the kernel with the knowledge of which version is making the request. This information can be passed, when necessary, to the system services level to allow differential handling, a procedure used in setpgrp, kill, and signal system calls.

System V and 4.2 BSD differ slightly in several system header files such as acct.h and sgtty.h. However, in most cases, the versions have a common set of structure members with the same names, each having a few unique members for other functions. By using a system header that is a superset of the common and unique members from each version, utilities and application programs running under 4.2 BSD or System V can run transparently to the existence of the other system. Typically, the amount of time required by the kernel to fill the superset structure compared to the time for the original structure is insignificant.

Some of the features of System V are unique enough that it was easiest to handle them with new modules incorporated into the system interface and system service levels. Examples include the System V semaphores, interprocess message facilities, and shared memory features. It would have been possible to put these on top of the 4.2 interprocess communication facility mechanisms, but this is inefficient and leaves the possibility of incompatibilities. The amount of code required to implement these features directly is relatively small, and the impact of interactions on the rest of the kernel is remarkably little. Overall, the OSx kernel is only nine percent larger than the standard 4.2 BSD kernel.

Difficult compatibility issues

Some other conflicts between different operating systems are more serious and require special mechanisms to handle them in a manner that maintains compatibility without sacrificing performance. These include handling the differences between System V's fixed-length file names in directories and 4.2 BSD's flexible-length approach; handling System V's unique named pipes, different signal handling procedures, System V interprocess messaging mechanisms; and handling terminal drivers and character I/O.

In approaching the differences in directory name structures, the significant difference from a system interface viewpoint is that 4.2 BSD uses a variable-length buffer to hold the character names of i-nodes. In contrast, System V expects the directory name buffer to be of fixed length. The 4.2 BSD's flexible length approach to directories is potentially of great advantage—especially when networked applications become more common. The problem is to provide a directory interface to System V utilities and applications compatible with what they expect, while still retaining a flexible-length directory structure underneath.

The 4.2 BSD uses a variable-length structure with an entry that defines the length of any given directory record. To resolve the conflict between the operating systems, OSx takes advantage of the fact that it knows the current process's universe. When a process issues a read system call, a flag is set in the process structure indicating that it is, for example, a System V read call. This flag is checked and, if it is a System V read, the variable-length directory structure read from the disk is converted to a fixed-length record, which is returned to the user.

For application programs, this approach is compatible with the original System V structure. The company has increased the fixed name limit within System V to be compatible with the maximum variable-length name in 4.2 BSD so that a System V user will never see a 4.2 BSD directory name truncated.

Other changes are of interest. For example, modifications were made to the stat system call so that System V programs, which compute the number of files in a directory by computing its size, do not need to be changed because of a flexible name format. The net effect on System V users is that they can think in fixed-length directory records while working in a variable-length directory file system.

Named pipes are a unique System V feature. With them, a process can set up a pipe to be used at some future time by a process that has no knowledge of the pipe's creator. Pipes are implemented in 4.2 BSD in terms of the more general socket interprocess communication mechanism. But, BSD does not implement this particular form of pipes.

In OSx, a new i-node type IFIFO was created (based on System V), and the 4.2 BSD socket mechanism was used to effect the actual piping. This was done in a similar manner to how unnamed pipes are currently implemented in 4.2 BSD. This design entails making changes in the code for opening, closing, reading, and writing from i-nodes, and utilizing the socket functions. A side effect of this implementation is that 4.2 BSD users can take advantage of named pipes if they so choose. However, in the interest of a clean 4.2 BSD interface, this feature is not recommended for 4.2 BSD programs.

Signal handling

The two operating system versions differ considerably in the way they perform signal processing. In terms of signal types, 4.2 BSD is nearly a superset of System V, and provides the starting point of OSx code. There are, however, significant differences. For one, signal handlers in 4.2 BSD are set and cleared through a library subroutine. This subroutine translates requests into sigvec calls, and saves signal masks. In System V, signal handlers are set through a signal system call.

In System V, a signal system call can be made with parameters specifying that when a signal signifying the termination of a process occurs, it can be ignored. If a wait is then executed, it will block until all of the processes created by the parent process have terminated. Moreover, when a signal is caught in 4.2 BSD, further signals are held or blocked. Neither option is available in System V.

OSx modifies 4.2 BSD to provide an additional signal system call. Upon entry into the kernel, a flag is set in the process structure for that process. It indicates that the process wants System V-based signal handling. Note that whether a process invokes System V or 4.2 BSD, signal handling is independent of the logged universe.

For interprocess communication, System V introduced three new mechanisms—semaphores, messages, and shared memory—along with the associated system calls. For efficiency's sake, and to avoid subtle incompatibilities, these features were incorporated directly into the kernel, rather than layered on top of 4.2 BSD sockets.

For semaphores and messages, the code ported almost without change from System V. The shared memory feature required porting to the virtual memory system, where hardware support within the System 90x was taken advantage of for shared memory pages. For shared memory, this hardware support is a new type of virtual segment specifically handled by the pager and swapper mechanisms.

Terminal drivers and I/O control

Providing dual 4.2 BSD and System V compatibility for the operating and control of terminal I/O was the single most difficult compatibility task in OSx implementation. Although the interface through the ioctl system call is essentially the same in both operating systems, the degree of control provided and the underlying implementations are radically different. In this case, neither operating system is a superset of the other, and each provides a subset of unique features as well as common ones.

The OSx supports both System V and 4.2 BSD views of terminal I/O. This allows users to switch back and forth between universes without leaving their terminal in unexpected states.

While basing the OSx driver on the 4.2 BSD terminal driver, the maze of conflicting, and partially agreeing I/O control call parameters were redefined to be non-conflicting, while not changing any of the symbolic names in either universe. This redefinition was done so that the terminal driver can detect, from a parameter, whether the calling process is 4.2 BSD or System V. In addition, the terminal state vector within the sgtty structure was modified to include a superset of the states known to the 4.2 BSD and System V universes. Finally, code was added to handle System V-unique I/O control parameters.

The kernel can detect which stty (set terminal state) system calls should be interpreted as transitions between universes, as opposed to settings within a single universe. To aid in doing this, calls are distinguished by a transparent 32-bit identifier that is added to the sgtty structure. This tag identifies the sgtty structure as being from a particular universe. The kernel then uses this information to determine if the terminal settings should be reset or default settings used when the stty call is made. In this way, the OSx operating system can support easy transitions between the two universes without resetting the user's terminal to an unexpected state.

The result of this software development for the OSx user is that programs that are based on either 4.2 BSD or System V terminal I/O conventions need not be modified to run under OSx. In addition, users can execute a combination of such programs within a single terminal session from either universe.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 731 *Average 732* *Low 733*

Join the power alliance.



Panasonic Lithium, Lead-Acid & Ni-Cd Batteries.

Meet your new allies - Panasonic Batteries, the "power elite." Batteries that are 100% tested, with the right chemistry and fit for your applications; a "power alliance."

Lithium Batteries: Extremely long life, very high energy density. Ideal for microprocessor memory retention. Available from stock; in pin, coin and cylindrical packages.

Nickel-Cadmium Batteries: State-Of-The-Art rechargeable performance. Cover wide range of portable and stand-by applications. Used extensively in instruments, telecommunications, medical equipment and computers.

Sealed Lead-Acid Batteries: Special design yields long life, and can operate in any position. Used world-wide for portable VCRs...computer memory... emergency lights...alarms and UPS systems.

With Panasonic Batteries, there is no compromise, no force fit. The world's broadest line...from one of the world's largest battery manufacturers. Batteries made to exacting Q.C. standards that place them in the top echelons for reliability.

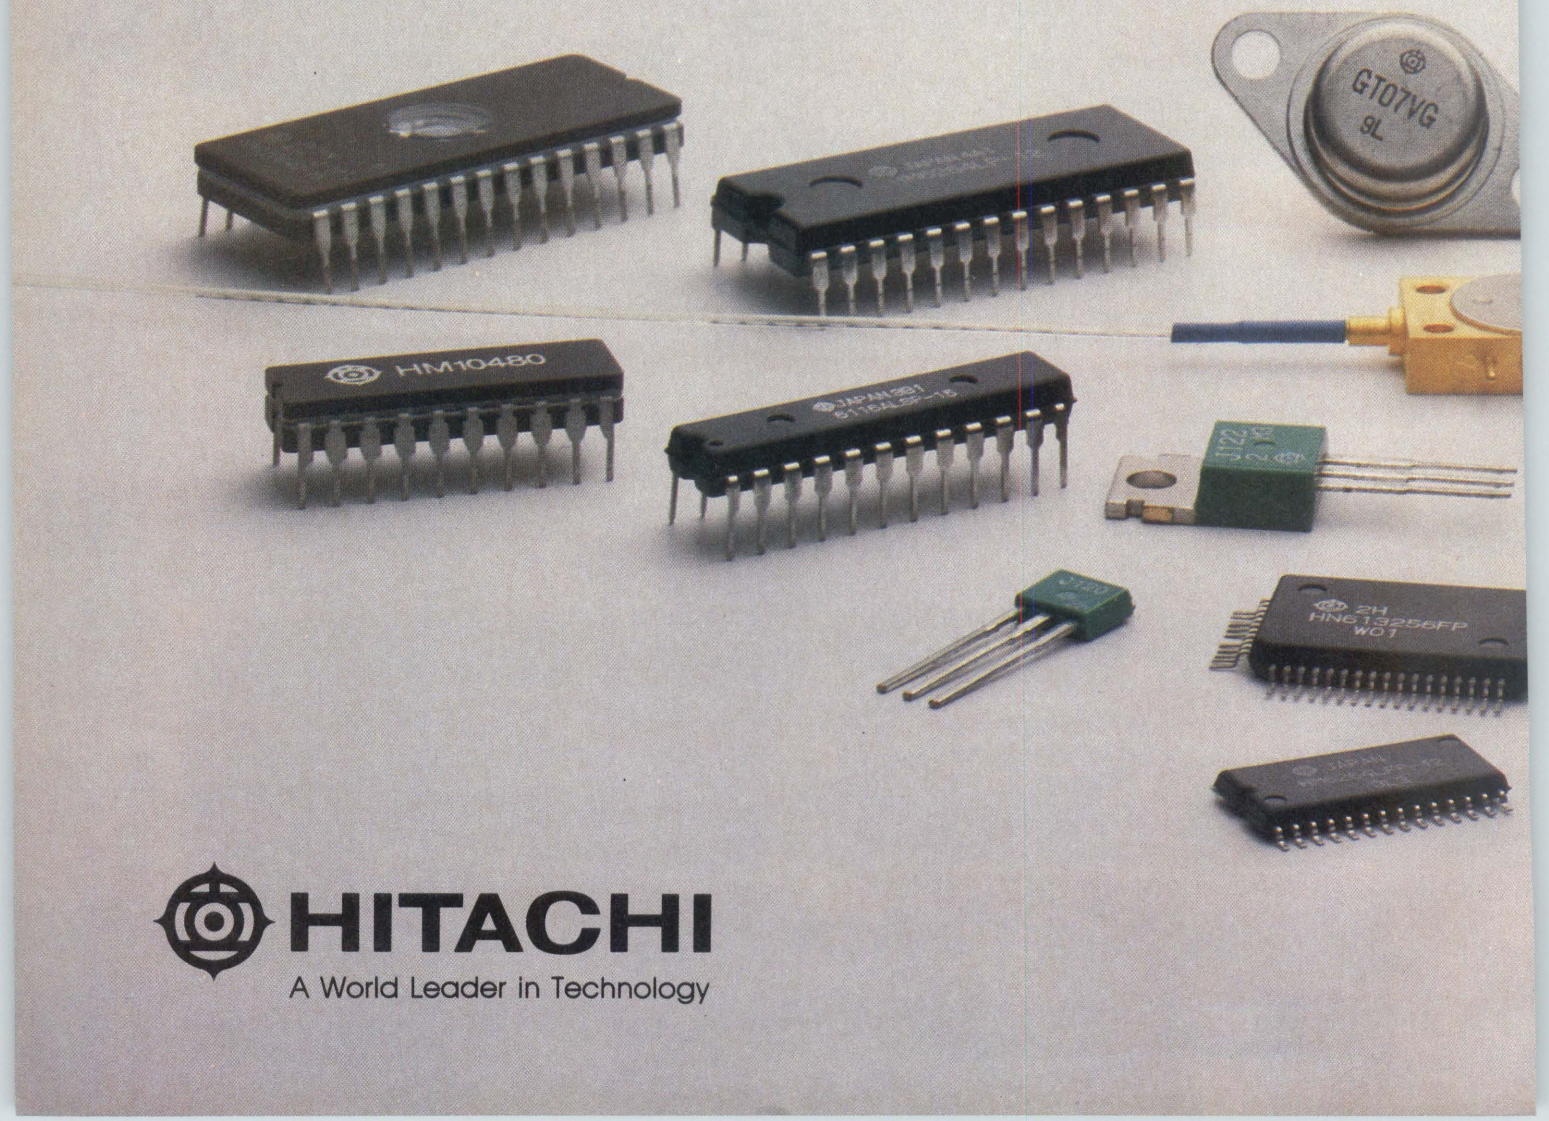
Join the "power alliance." Write or call: Battery Sales Division, Panasonic Industrial Company, Division of Matsushita Electric Corporation of America, P.O. Box 1511, Secaucus, N.J. 07094; (201) 348-5266.

Panasonic®
Industrial Company

CIRCLE 103

HITACHI

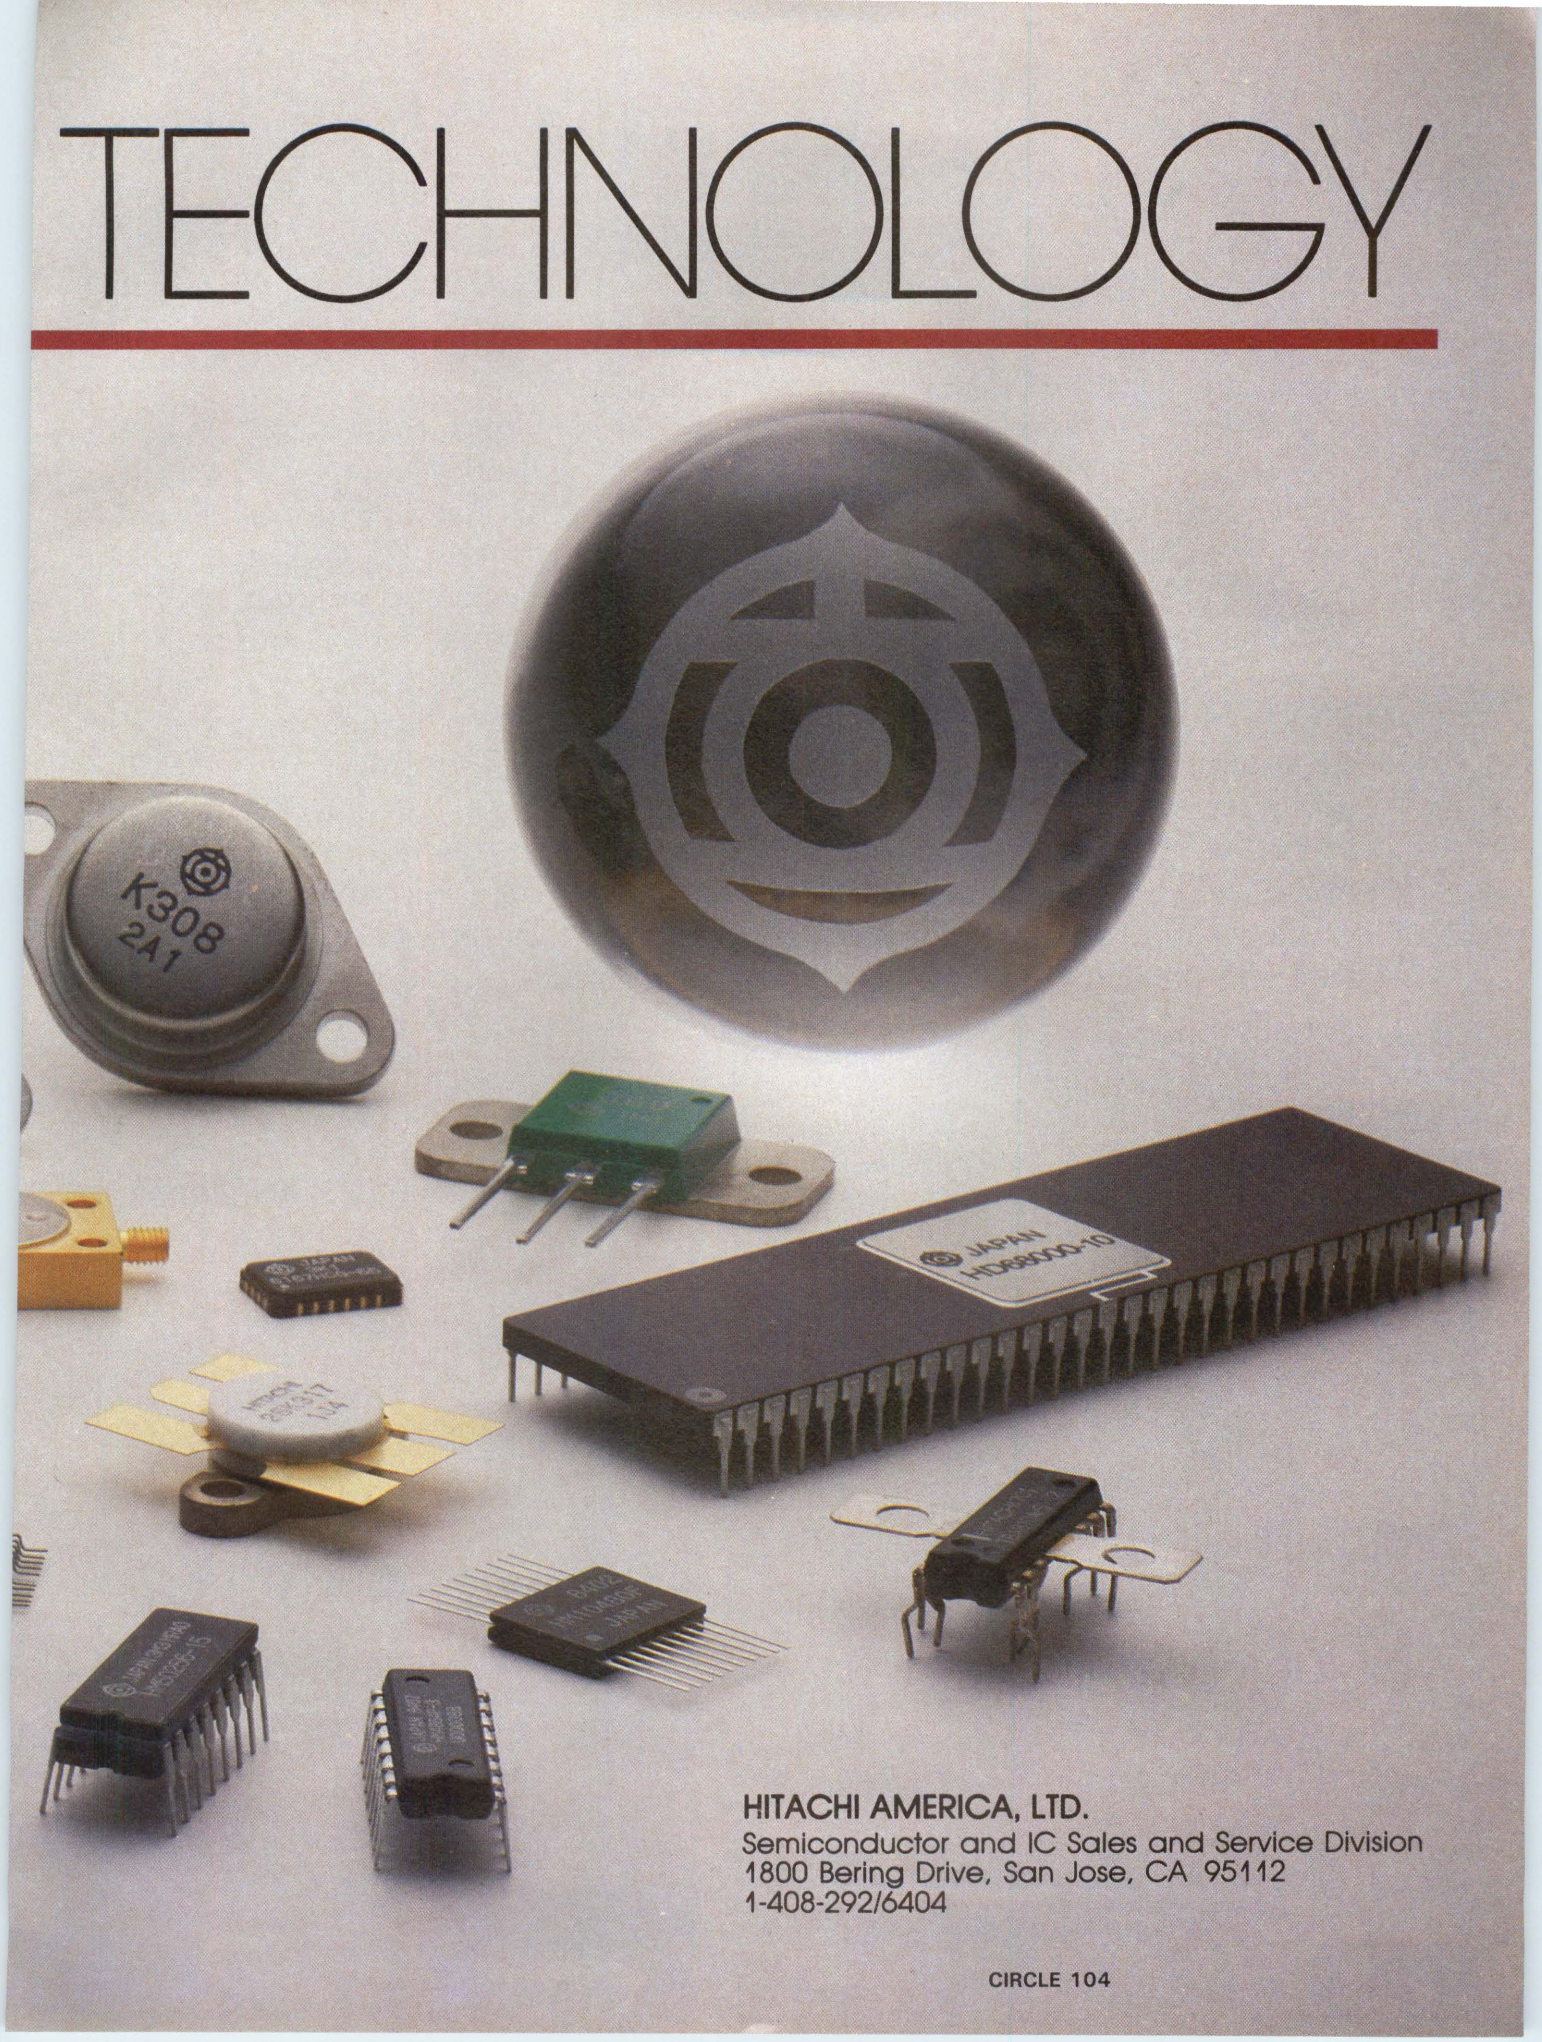
Expanding the Limits of Excellence in Semiconductor Design



HITACHI

A World Leader in Technology

TECHNOLOGY

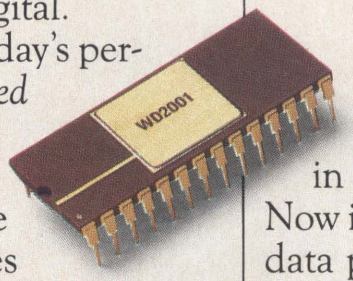


HITACHI AMERICA, LTD.
Semiconductor and IC Sales and Service Division
1800 Bering Drive, San Jose, CA 95112
1-408-292/6404

CIRCLE 104

How to build a really

Put the "personal" back in personal computers with the WD2001 Personal Chip™ from Western Digital.



Consider the facts: today's personal computer is a *shared* computer. People in the Modern Office are standing in line to use the nearest PC. John updates his sales forecast, Jane keeps her product plans on target, Martha processes salary reviews, Joe does his monthly report, Teddy wants to write a letter home to Mom—all on the same PC.

How can they keep their files private and prevent electronic eavesdropping? With a Personal Chip, that's how.

Your customers will take it personally.

In this Age of Compatibility, it's tough for a PC maker to build in distinguishing product features. Now it's easy to add a unique feature—data privacy—and offer your customers a more personal computer. Just design in the Personal Chip.

The WD2001 Personal Chip is a complex-on-the-inside, simple-on-the-outside VLSI device that enables users to lock their files. It implements a sophisticated encryption algorithm that's been exhaustively tested and certified by the National Bureau of Standards.

Each user has a personal key, known only to him or her. There's not even a look-up table



personal computer.

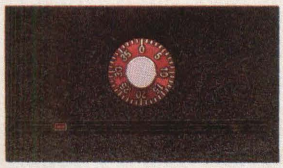
for clever hackers to stumble upon. Without the key, stored data is unintelligible and uncrackable. Type in your key and, instantly, alphabet soup is turned back into data.



For your eyes only.

We humans are naturally nosy. We can't resist a peek at our neighbor's mail if we get the chance. With a shared computer you're just tempting human nature.

Winchester disk drives in personal computers compound the problem of data privacy. At least you can lock your floppy disks in a drawer, or even take them home. If your data is on a hard disk it is there for everyone to access.



Put a lock on data stored on disk with the Personal Chip.

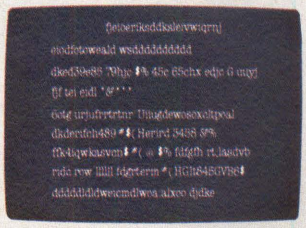
The Personal Chip protects data privacy on floppies, hard disks, tape drives, shared peripherals, even during transmission in a local area network. In fact our NetSource/PC-LAN IBM-compatible local area network board includes a Personal Chip to ensure secure data transmission.

Our application engineering team is ready, able and eager to work with you to help you design in the Personal Chip. Because the WD2001 takes just a single socket, it makes sense to design it into your system's motherboard. Or, as the industry's leading designer and manufacturer of disk controller boards, we're in a perfect position to supply you with a custom Winchester/floppy/tape controller with a built in Personal Chip for secure data storage.

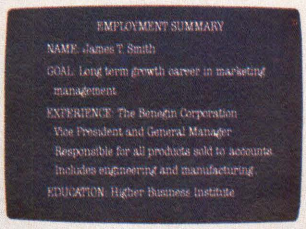
Personal Chip is a trademark of Western Digital Corporation.

Your technology partners in data privacy.

Make us your technology partner



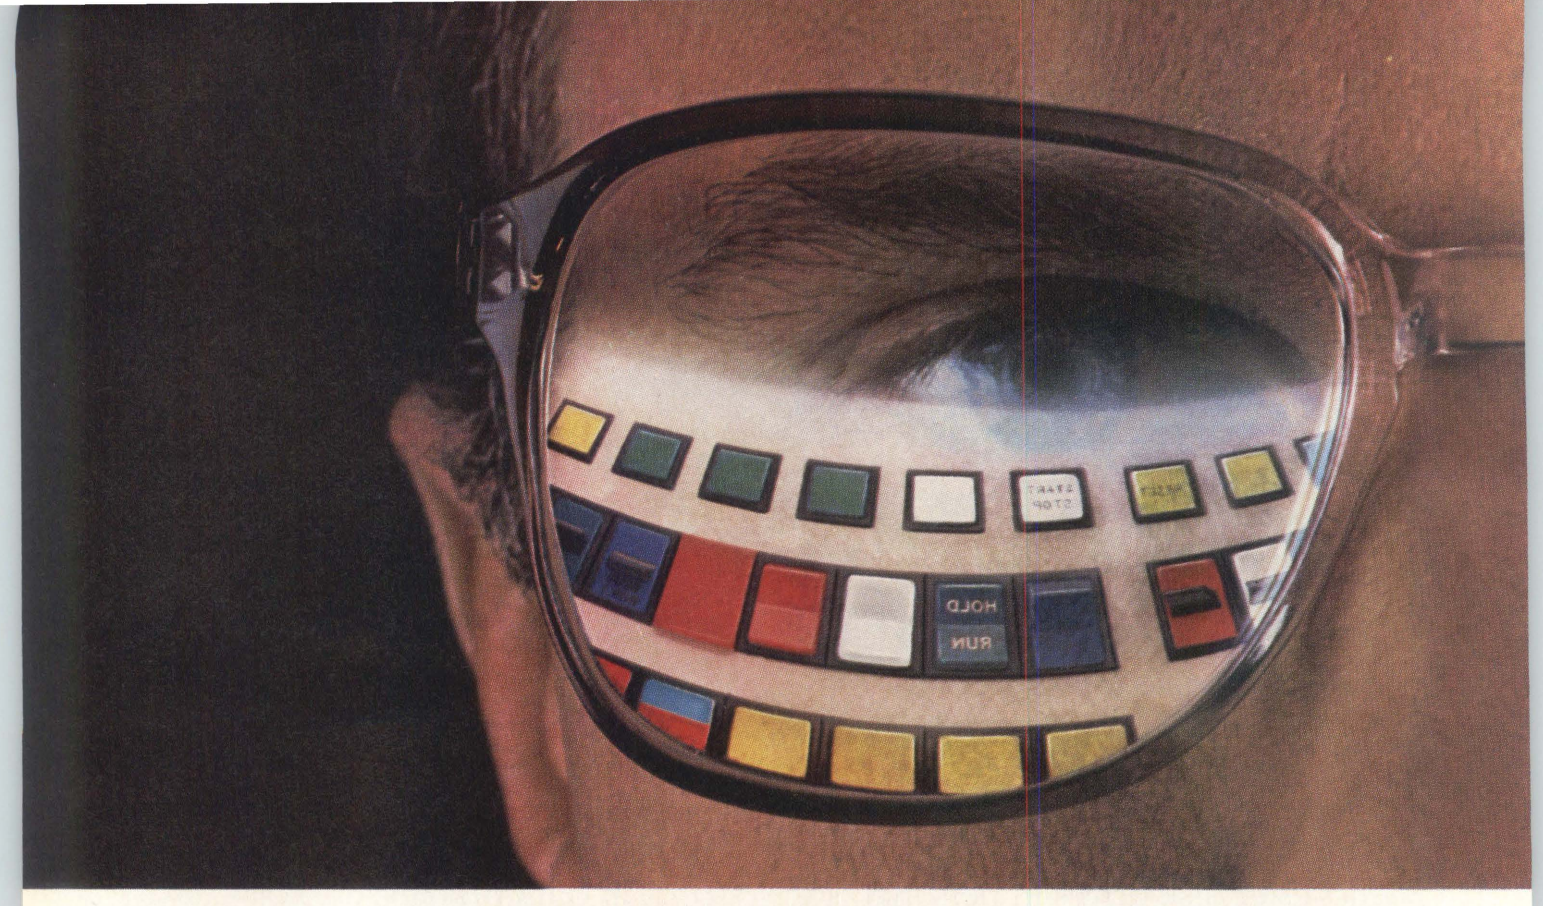
The Personal Chip encrypts data, making it safe to store or transmit.



Decrypt data with a few quick keystrokes—if you have the key.

in data privacy and we'll work with you closely to ensure your success. Write for the WD2001 Data Sheet and Application Note. Or give Dan Fili a call on our Data Privacy Hotline, 714/863-7828. Western Digital Corporation, 2445 McCabe Way, Irvine, California 92714.

WESTERN DIGITAL
CORPORATION



We design our through your

With the first glance, your customer begins to form an impression of your product. And the manual controls you use play a major role in this important first impression—as well as long-run operator satisfaction.

That's why we take great care to build human factors into the design of our manual products. By designing our manual controls from your customer's point of view, we can help make your product look better, feel better, work better and, finally, sell better.

Not only do you get well designed manual controls from MICRO SWITCH, you get lots of them. Our continuing development program has led to a variety of new products to fit specialized requirements. And with the broadest selection of product options in the industry to choose from, you get the flexibility and freedom to design your products to meet your customer's needs.

New looks for new needs.

Our Programmable Display Pushbutton is an excellent example of our ability to respond to emerging needs in the marketplace with new products having unique human factors features. It is a user-friendly pushbutton that provides two-way communication between the system and the operator. Our Programmable Display Pushbutton offers the state-of-the-art ability to display an infinite number of legends when programmed through a computer. This multi-function capability means that one switch can replace dozens, reducing panel space and viewing area to simplify the operator's tasks and reduce the possibility of error.



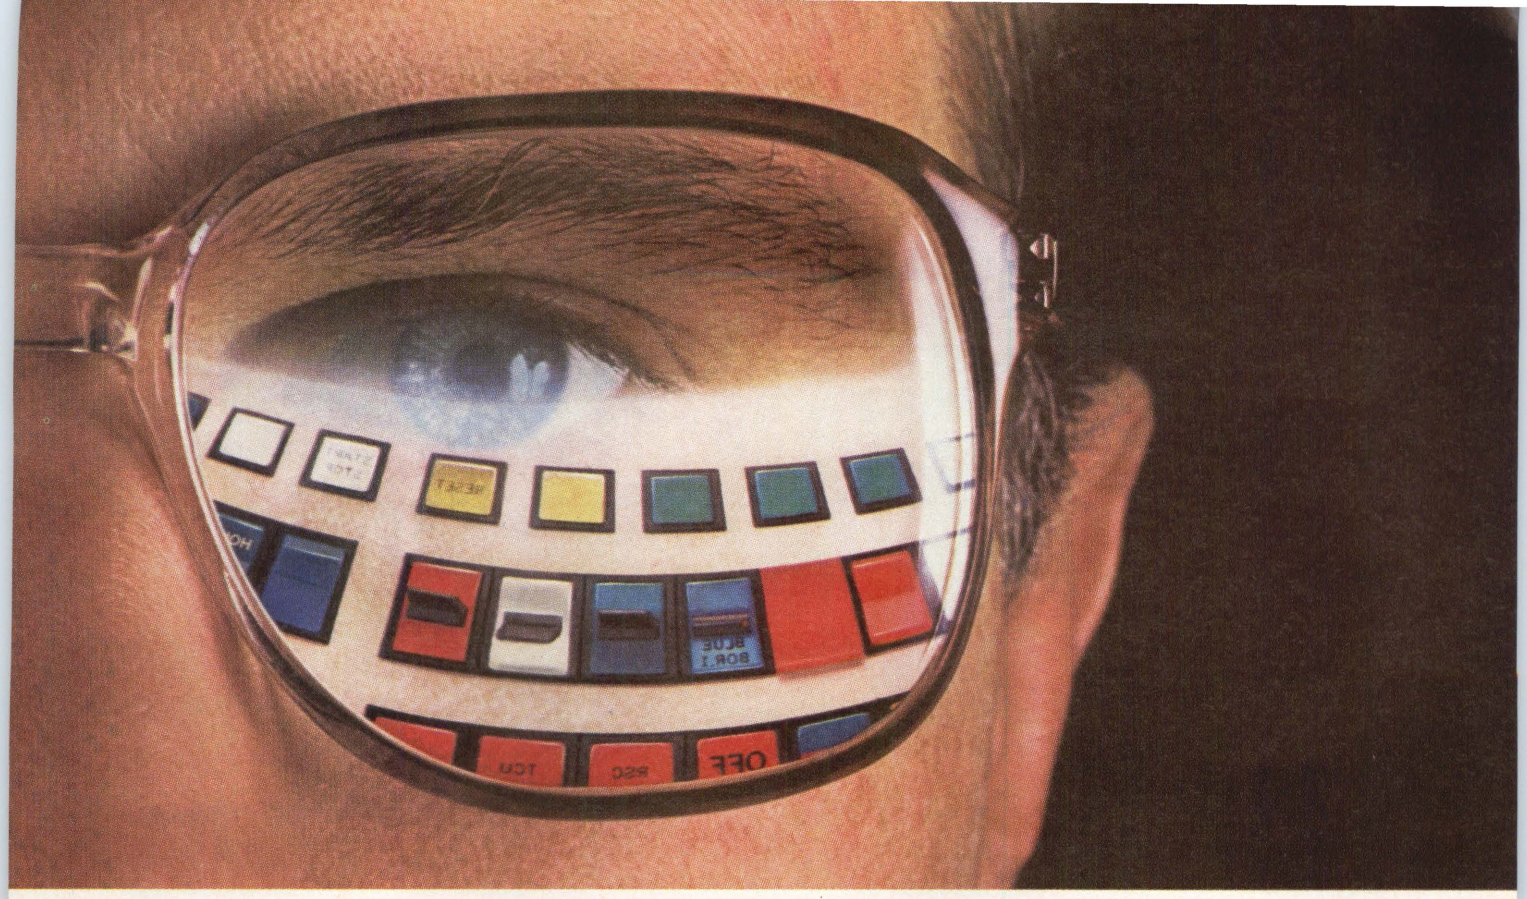
Even the best can get better.

We're always working to improve our manual products to meet your changing needs. For example, based on years of experience in building quality keyboards, we were able to combine quiet operation and excellent tactile feedback in our new Silent-Tactile keyboards. These low profile keyboards look good and feel good. They retain their consistent positive feedback month after month, year after year. Their silent operation helps make work environments more pleasant. Because in a room full of CRTs, it's what you don't hear that counts.



Take your pick from a variety of quality controls.

Constant efforts over the years have resulted in a broad line of human factors-designed manual products. Our touch panels, for example, combine high quality graphics with reliable switch operation. We take special care with such details as embossing for positive finger positioning, snap discs for tactile feedback and high quality color and graphics.



manual controls customer's eyes.

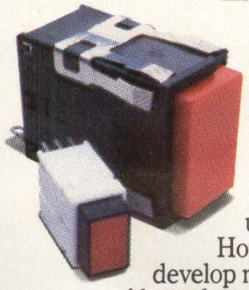
The result: our products help make your product easy to use and good looking, as well as reliable.

The Advanced Manual Line and Miniature Manual Line display the same painstaking concern for detail and quality. They're designed to be visually pleasing—with clean lines, vivid colors, bright, even illumination, attractive shapes and sizes. And their appearance is coordinated, both within and among product groups. So you can build harmony into your design.

Having this wide selection of controls to choose from means you can specify the best combination of performance and price—always with superior human factors features.



Experience you can count on.



For over 40 years we've been building manual products to help solve all kinds of control problems. For years we've been conducting and publishing the results of original human factors research. No other controls manufacturer does. And we work with Honeywell's industrial design experts to develop new products and materials.

Along the way, we developed special capabilities in plastics, inks, illumination and switching techniques that we can apply to all our manual products.

Human factors design ideas—yours for the asking.

We have compiled a wealth of basic human factors design considerations in a 36-page booklet—our Panel Design Guide. You'll find it an invaluable reference tool in planning and designing the interface between man and machine. Simply call or write to get your copy.

The response and support you need.

Whether you need a standard product or a custom solution, you can count on MICRO SWITCH for the help you need. Our experienced field and application engineers are ready to work with you to find the best solution for your project.

No matter where you are, MICRO SWITCH is ready. For local availability of standard products, just call on one of our Authorized Distributors. And there's our broad network of sales and service representatives around the world.

For the location of a sales office or authorized distributor near you, call 815-235-6600. Or, write MICRO SWITCH, Freeport, IL 61032.

We'll help you find the manual controls and human factors features that work best for you and your product—no matter which way you look at it.

Together, we can find the answers.

MICRO SWITCH
a Honeywell Division

CIRCLE 106

Announcing the new generation in eight-inch, high-capacity disk storage.

Enhance your system's performance with the 300-megabyte storage and ultra-fast response of the all-new DX-300. Sized to fit in an eight-inch floppy drive envelope, it meets the demands of multi-user, multi-access systems for increased database, fast throughput and, of course, unparalleled reliability.

The DX-Series is based on the *field-proven* technology of standard oxide media and reliable composite heads. We've added a completely new voice coil positioner and some very innovative engineering in both mechanical design and electronics.

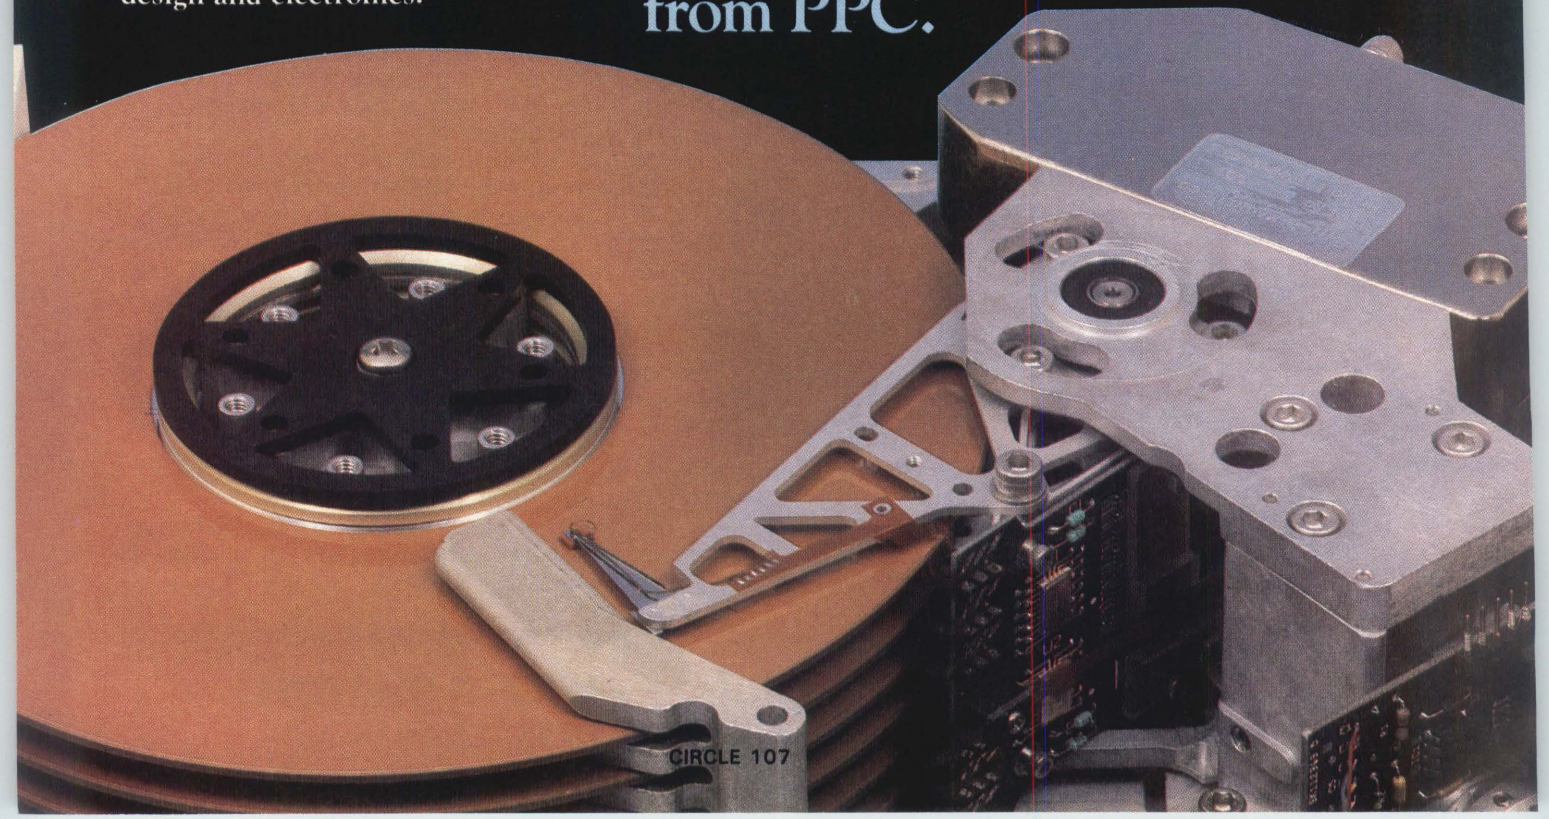
The DX-Series offers capacities of 180, 240 and 300 megabytes, with ANSI and SMD interfaces to quickly and easily integrate its high performance characteristics into your system.

Move into the high end of cost-effective data storage with the *proven* performance and reliability of the DX-300. Write or call for complete technical data and fast and firm delivery dates.

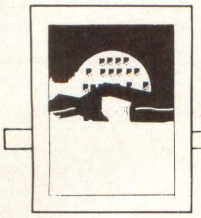
The DX-300, 300-megabyte Winchester disk drive from PPC.

TA PPC PERTEC PERIPHERALS CORPORATION

9600 Irondale Ave. Chatsworth, CA 91311
(818) 882-0030 TWX: (910) 494-2093



CIRCLE 107



MINICOMPUTER SYSTEM OFFERS TIME-SHARING AND REALTIME TASKS

An adaptive job scheduling algorithm, and both data and directory caching help optimize minicomputer operating system performance.

by Phil Sherrod and Stephen Brenner

Minicomputer operating systems usually perform certain tasks well at the expense or exclusion of others. For example, single-user operating systems, which are common on machines with small amounts of memory, provide excellent response time and hardware access. In contrast, multi-user operating systems often sacrifice response time or hardware access to handle more than one customer.

Similarly, a realtime operating system is usually not very good at time-sharing. Unfortunately, the ability to control minicomputer hardware is often critical to the performance of realtime tasks such as instrument control or data acquisition. One operating system that attempts to solve this problem is the TSX-Plus. This operating system's objective is to provide a multi-user minicomputer operating system with good time-sharing response time. It also aims to concurrently give privileged jobs the hardware access essential for realtime tasks.

The operating system can be used by computer designers and system integrators for Digital Equipment Corp's LSI-11 and PDP-11 minicomputers. These are designed to provide facilities concurrent with time-sharing operations. TSX-Plus is compatible with RT-11, DEC's most widely used single-user operating system. It extends, and is used with, RT-11 to offer a time-sharing environment for RT-11's command structure and programming interface.

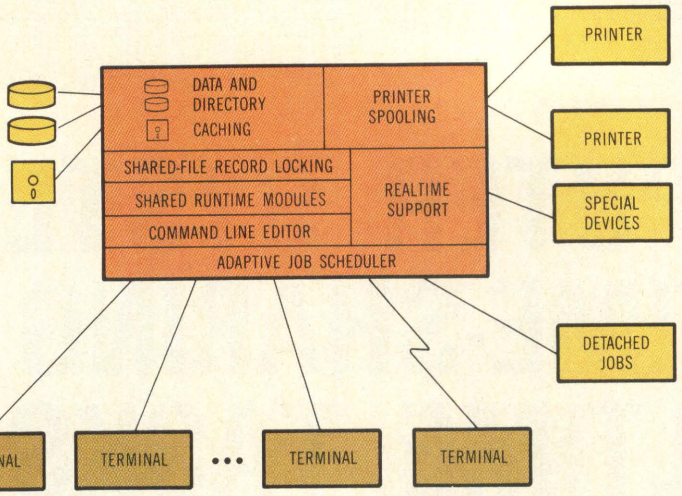
When started, the operating system replaces the RT-11 monitor kernel and separately provides its system service calls. As a result, it runs most RT-11 programs without recompilation. In addition, it features virtual task windowing through virtual lines and detached jobs, an adaptive job scheduler, and high speed interrupt servicing (see the Figure).

Phil Sherrod is vice president of software development at S&H Computer Systems, Inc, 1027 17th Ave S, Nashville, TN 37212. Mr Sherrod holds a BA in physics from Vanderbilt University, Nashville, Tenn.

Stephen Brenner is senior systems programmer at S&H Computer Systems, Inc. He holds a BS in chemistry from the South Dakota School of Mines and Technology, Rapid City, SD, and a PhD in biochemistry from the University of Minnesota, Minneapolis, Minn.

Job swapping

To efficiently use TSX-Plus as a design tool, it is necessary to understand how it handles time-sharing and real time with the same software. It evolved from the TSX operating system, which was developed as an extension of RT-11. TSX ran in conjunction with RT-11 to provide small, inexpensive DEC-like computers with time-sharing services but without the use of memory management hardware. In contrast, other time-sharing operating systems for PDP-11s required memory management hardware



The operating system provides multi-user time-sharing concurrently with realtime support. It optimizes system performance with an adaptive job scheduler, and data and directory caching. With it, the computer designer has virtual task windows, detached jobs, shared-file record locking, and command line editing.

that was only available on relatively expensive computers. TSX provided its time-sharing services by extensive job swapping in the limited memory space (less than 56 Kbytes) available without memory management. In addition, it provided such features as transparent printer spooling and shared-file access control.

As the PDP-11 line evolved, especially with the development of the LSI-11/23 processor, relatively inexpensive computers became available with the memory management hardware necessary to address more than the 64-Kbyte virtual address space of the PDP-11 series. Although still providing time-sharing services and realtime control, TSX-Plus has been expanded to take advantage of the increased memory on newer computers. To do this, it uses an adaptive job scheduling algorithm to provide rapid time-sharing response while maintaining the flexibility to promptly service realtime interrupts. All this is possible with a further benefit to the computer designer: the operating system is set up so that the broad range of RT-11 application software is compatible with it. This feature provides the system integrator with an easy upgrade path for current RT-11 users.

Most time-sharing operating systems available for the PDP-11 series sacrifice power, ease-of-use, real-time services, speed, flexibility, or size to run the machines. In fact, of all operating systems available for the PDP-11 series—including RT-11, RSX, RSTS, V7M-11 (Unix), and UCSD p-System—only TSX-Plus provides multi-user, multitasking, time-sharing, rapid response, fast execution, and realtime services in a memory-efficient environment.

There are other TSX-Plus practical features of interest to the computer designer. For example, it supports the full addressing range of LSI-11 machines, which address 4 Mbytes of memory. Similar extended addressing support is available for the PDP-11 (Unibus) machines even though RT-11 does not yet support extended Unibus addressing. Also, the latest version of the operating system supports virtual extended memory addressing (including virtual arrays and virtual overlays). This is available through Fortran and other compilers. With this feature, individual programs can exceed the 64-Kbyte virtual address limitation that is inherent in the PDP-11 hardware architecture.

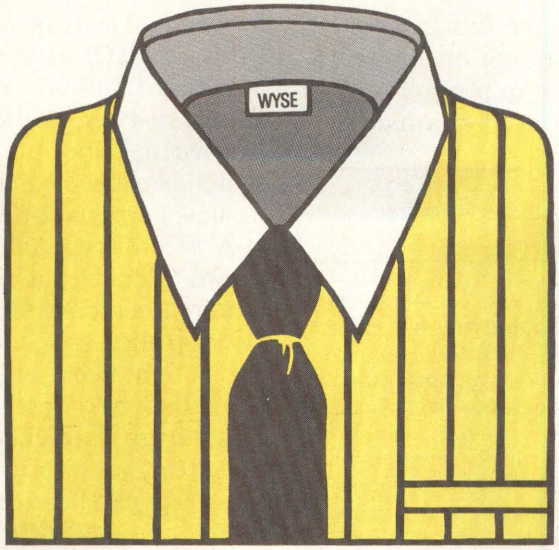
The latest word

Version 5.0 of TSX-Plus, released in April 1984, includes the adaptive scheduling algorithm (see Panel, "How an adaptive scheduler works"). The algorithm, which is responsive to time-sharing jobs, allows time-critical tasks to execute on demand and low priority jobs to utilize idle minicomputer time. Moreover, a software mechanism is implemented to permit even faster program inline interrupt servicing. And, a generalized data caching facility is added to maximize system throughput. For older machines that have been upgraded to 22-bit addressing (up to 4 Mbytes of memory), a software facility is included to permit the use of existing older DMA devices (ie, disk drives) that were designed for 18-bit addressing (256 Kbytes).

The operating system provides virtual task windowing—one user can control more than one process (see Panel, "Using virtual task windows"). Rather than simply allocating a portion of the screen to an alternate process, the operating system allows users to switch among two or more entirely separate tasks from the same terminal. This can be done at any point in a program or at command level with a simple two-character command.

The execution of the original task may continue (although at a lower priority) even when it is not connected to the terminal. For example, it might be

Announcing the WY-75.



Our new WY-75, VT-100*software-compatible terminal is specially fitted for use with your DEC*computer.

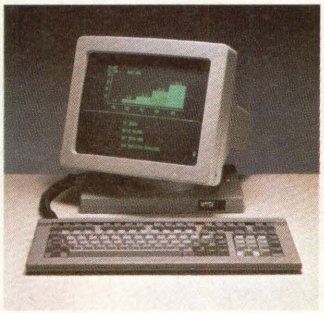
It offers a combination of features you can't find in any other VT-100 software-compatible terminal. Like a compact, ergonomic design. A finely sculpted, low-profile keyboard. And a swivel and tilt non-glare 14" screen, tailored with an 80/132 column format.

At a price of $795, the WY-75 won't cost the shirt off your back.

Contact Wyse Technology for more information. And put your DEC on the best dressed list.

$795

VT-100 software compatible and dressed for success.



WYSE

Make the Wyse Decision.

Wyse Technology, 3040 N. First Street, San Jose, CA, 95134, 408/946-3075, TLX 910-338-2251, Outside CA call toll free, 800/421-1058, in So. CA 213/340-2013.

*VT-100 and DEC are trademarks of Digital Equipment Corporation.

CIRCLE 108

very convenient to start a sort operation as one job and then perform some text editing or data entry as a second job while the sort is in progress. In an ordinary time-sharing environment, performing such tasks might only be possible by first submitting the sort as a batch process, or alternatively, by using a second terminal—if a free terminal is available. With TSX-Plus, one need only start the sort, switch to a virtual line, perform any desired operations, switch back to the previous job, and continue from its next terminal operation.

A second type of multitasking is the ability to spawn "detached jobs." These are tasks that do not require interactive terminal operations and that can be controlled by a command file. With this operating system, such jobs can be initiated from within programs, by a keyboard command, or during system startup. Detached jobs can be used as common resources, monitoring programs, or for any other tasks that can be executed independent of operator intervention. TSX-Plus provides a system of "message channels" through which tasks can communicate with each other and with detached jobs.

One of the most noticeable features of a time-sharing operating system is its response to an interactive terminal operator. The operating system's job scheduling algorithm gives priority to terminal operations. Thus, when an operator enters data or a command, the execution of any lower priority task is interrupted and the new terminal input is quickly processed. If a job executes for a significant amount of time without terminal interaction, then it is assigned a lower priority. This procedure gives preferential treatment to more interactive jobs. The process continues, frequently reassessing the "interactive" nature of each job, and gives the best response to those jobs that are most heavily operator dependent.

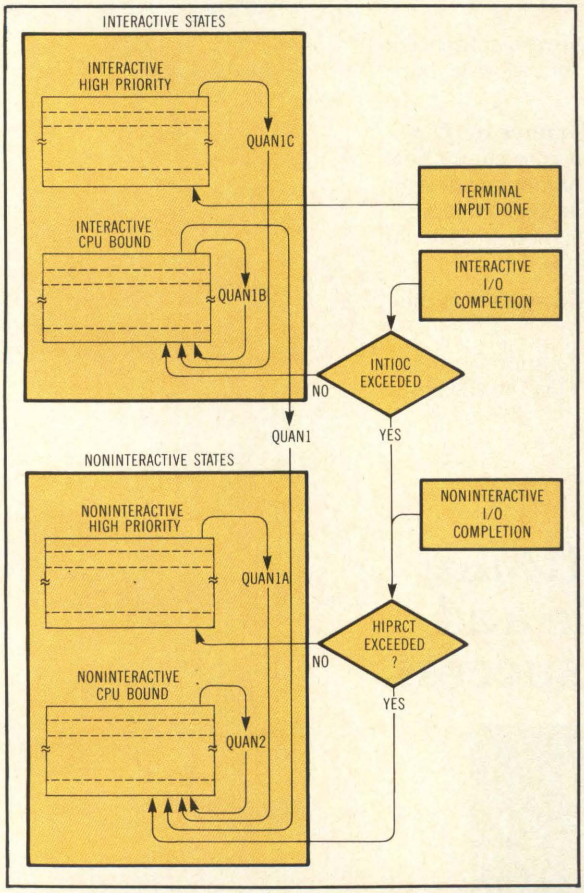
If a job is waiting for an operation to complete (such as a timed event or an I/O operation), the next available job is immediately scheduled. The rapid response permitted by this method gives precedence to those jobs that are most operator intensive, while at the same time giving time to low priority jobs whenever it is available.

The adaptive scheduler allows for better performance and greater flexibility than operating systems that use simple round-robin scheduling. With round-robin scheduling, every job gets an equal time-slice, regardless of its priority. The adaptive scheduler is also superior to a full priority-driven scheme that cannot dynamically adjust itself to changing job characteristics.

While adaptive job scheduling best serves the normal time-sharing operation, some circumstances require special consideration. Jobs that control instruments or that perform data acquisition often need immediate response. On the other hand, it might be desirable to allow some very low priority jobs to execute only when the system is not busy with any other job. The operating system provides job priority ranges both above and below normal interactive jobs to handle such special cases.

How an adaptive scheduler works

The adaptive job scheduler optimizes terminal response for interactive jobs. Although very high and very low priority jobs are possible, most time-sharing jobs are scheduled according to their activity. When a job completes a terminal input operation, it is placed in the highest interactive priority queue (see the Figure). Each job in this queue is run for a very brief period and then rescheduled in a slightly lower queue. Jobs in this next queue are each run for a slightly longer time before being placed in a still lower queue.



Whenever another terminal input is completed, the job is returned to the highest queue. If a job is waiting for a read or write to disk, it becomes non-executable. When the disk I/O completes, the job is returned to an executable queue determined by the number of read/write operations and the accumulated time. By separating time-sharing jobs according to their execution state and maintaining both execution time and I/O counters for each job, TSX-Plus is both able to provide rapid terminal response and still distribute system resources equitably among all jobs. The system may be fine-tuned by adjusting the values of various time and I/O count limits either during system generation or with keyboard commands.

In the high priority range, jobs are executed strictly according to their assigned priority. When such jobs enter a wait state, lower priority jobs, such as interactive time-sharing jobs, are serviced. In the very low priority range, jobs are also executed strictly according to priority. But, since their priority is lower than interactive jobs, they execute only when no high priority or interactive jobs need attention.

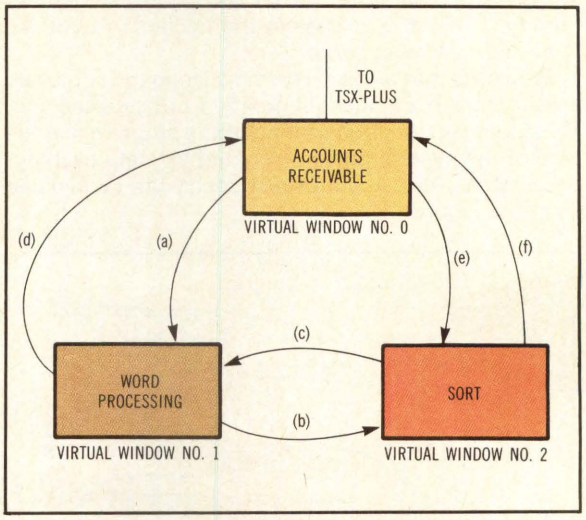
TSX-Plus uses many other optimization techniques in addition to the adaptive algorithm. One good example is illustrated with a data caching application (see Panel, "Data caching increases system throughput"). Data caching is a technique for keeping some of the data read from mass storage devices in memory. If the same information is needed again later, it is automatically supplied from the memory cache without rereading the disk. Since moving data around in memory is usually much faster than loading it in from disk, this technique can greatly improve the performance of programs that frequently reread external data.

TSX-Plus provides both generalized and shared-file data caching. Generalized data caching is effective when 50 Kbytes of memory or more can be allocated to the caching buffers and are applied to all data from all mounted devices. Shared-file data caching is most effective when only a small amount of memory can be devoted to data caching (less than 10 Kbytes), and is only applied to files that have been specifically declared as "shared files." Shared-file data caching is very effective with indexed sequential access method (ISAM) files.

To reduce long directory searches, file directory entries for all mounted devices are also cached. Both data and directory caches use a "write-through" technique to ensure data integrity in the event of a hardware failure. Any time cached data is modified, it is changed both in the cache and on the disk. When the caches become full, old information is replaced as necessary so that active files will benefit most from caching. Information brought into the caches by one user is available to all. In this way, commonly accessed programs such as editors and other utilities are likely to be available in the cache.

Using virtual task windows

Each operator can control several jobs simultaneously by switching between virtual task windows. A virtual task may be entered at any time (even while executing a program), by typing a simple two-character sequence. A virtual task not currently attached to a terminal continues to execute until it needs further terminal input or until its terminal output buffer becomes full. To understand this procedure, consider a business application, as shown in the Figure.



During the daily accounts receivable posting, a manager needs a letter keyboarded that must go out immediately. In the Figure (a), the operator, not wanting to interrupt the posting run, switches to virtual window No. 1 to create the letter. Unfortunately, before finishing the letter, the manager also requests a customer list sorted by zip code as soon as possible. In (b), the operator switches to another virtual window and initiates the sort. In (c), while the sort proceeds, the operator goes back to finish the letter, starts it printing, and logs off that virtual line (d). Logging off automatically returns the accounting program. When the sort finishes (e), the operator switches back to virtual line No. 2, prints the sorted list, and returns to finish the posting (f).

Getting service

While common external devices such as disk drives and printers are serviced by device drivers (handlers), less common devices found in laboratories may not have readily available device handlers. Realtime facilities provide support for user-written programs to service unusual devices. This is possible without the need for the programmer to learn the arcane techniques needed to write a special-purpose device handler. For example, in the PDP-11 architecture, devices are serviced by control registers mapped into a special range of memory locations called the "I/O page." Addresses in the I/O page are not normally available to time-sharing programs, but by using realtime services, privileged programs may access these special addresses, or more generally, may access any memory location. Realtime jobs may also control their own priority, prevent themselves from being swapped, or may even disable time-sharing altogether when necessary for critical processing.

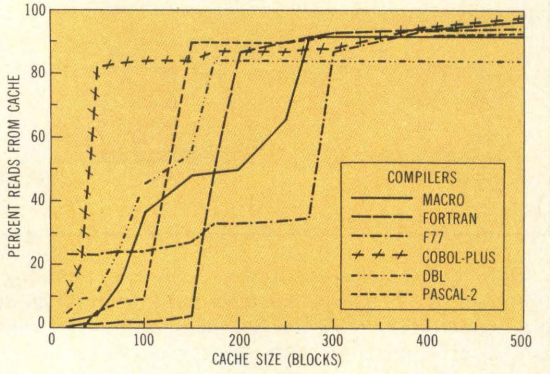
The PDP-11 architecture allows external devices to asynchronously demand system attention through vectored interrupts. Each vector normally points to the address of a device-dependent handler to service the interrupt. TSX-Plus allows time-sharing programs to service vectored interrupts by two methods. Inline interrupt service routines connect time-sharing programs to vectored interrupts with an absolute

Data caching increases system throughput

Data caching is a technique for maintaining copies of disk blocks in a system memory buffer. When any program requests data already in the buffer, the system automatically provides it from the buffer without performing another disk read. Since memory-to-memory data transfers are usually much faster than disk-to-memory transfers, this procedure can dramatically improve system throughput.

The data caching operation is completely user transparent. To avoid the possibility of data corruption in the event of a hardware failure, the caches are "write-through." That is, if data in the cache is rewritten, then the disk is updated also.

Programs that are overlayed, as language compilers frequently are, can benefit greatly from data caching. Depending on the organization of the program and the size of the system buffer allocated for data caching, as much as 95 percent of disk reads can be avoided (see the Figure).



minimum of system overhead, and can service interrupts at relatively high rates. Alternatively, interrupt completion routines are relatively free of restrictions and can service interrupts that are less time sensitive; these jobs may even be swapped out of memory when they are not servicing interrupts. To effectively shut off special devices when they are not in use, realtime programs may specify a list of devices to be reset when the program is terminated or aborted.

In addition to the special functions described, the operating system provides a range of time-sharing services such as printer spooling, shared-file access control, and interjob communication through named message channels. It also allows flexible job memory environments, shared runtime systems that may be mapped into several jobs' virtual image to conserve total memory usage, system and file access controls, indirect command files which accept parameters, user-definable commands, and a modifiable user command interface. The computer system designer also enjoys flexible terminal control from within programs, program performance monitoring, and simplified program debugging. As might be expected, to conserve memory space, the operating system supports optional inclusion of some of these features.

Applications galore

The computer system designer may consider TSX-Plus for general time-sharing use for a business environment, program development, and realtime instrument control and data acquisition. As an operating system suitable for business needs, it provides the rapid response necessary in a busy office. Because office staff are frequently not technically oriented, the command structure and the ability to define special purpose commands or to altogether replace the command structure with a menu are valuable aids in making the system easy to use. To protect important data, the operating system also provides shared-file record locking and robust file and directory handling.

System response is also very important when application programs are developed. Several screen editors are available for TSX-Plus along with a full complement of programming languages such as Cobol, Dibol, Basic, Fortran, Pascal, C, Forth, and a macro assembler. As previously stated, productivity is enhanced through virtual task windowing, which permits the programmer to perform multiple tasks simultaneously. TSX-Plus includes an integral program debugger that can be invoked at any time and does not have to be linked with the program being debugged. There is also a performance analysis monitor that will produce a histogram showing the execution time spent in each program section.

Instrument control and process control or data acquisition often require the ability to control peripheral interface hardware. TSX-Plus provides access to the I/O page and the means to service device interrupts from user programs with minimal system overhead. For high speed or other critical applications, programs may also disable time-sharing or otherwise control system functions such as job priority and swapping.

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 734 *Average 735* *Low 736*

Multibus® Microcomputer Systems Made Easy

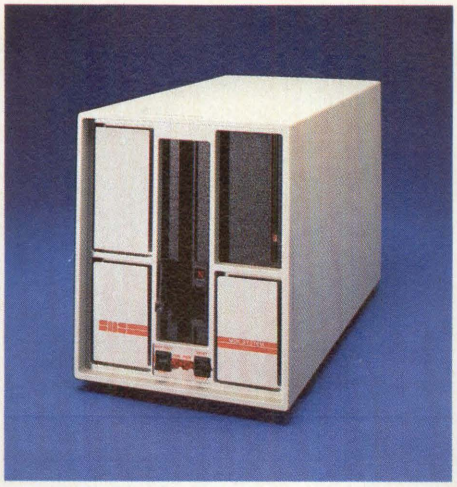
OEMs get products to market faster with SMS Winchester controllers, storage systems and microcomputer foundation systems.

Whether you select 8086, M68000, Z8000 or another Multibus single board computer, you can choose the SMS product to fit your packaging needs.



For large configurations, choose the DSX80000 foundation system with eight Multibus slots, 10, 20, 40 or 80Mb of Winchester storage plus an 8" IBM compatible floppy.

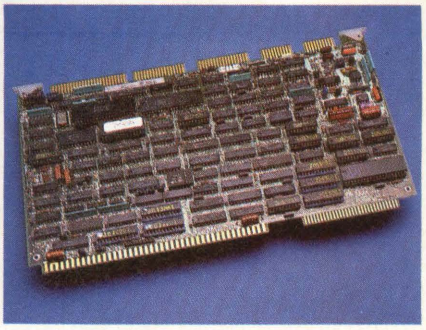
Small table top applications can use the MDX80000 foundation system with five Multibus slots, 10, 15, or 40Mb of 5-1/4" Winchester storage plus an 8" IBM compatible floppy.



If you have your own Multibus backplane, plug in the SMS FWT80000 storage system. It occupies only 5-1/4" of rack space and comes with either 10, 20, 40 or 80Mb Winchester storage plus an 8" IBM compatible floppy.



For high volume or unique packaging requirements, the SMS FWD8001/FWD8006 single board controllers are INTEL 215 compatible and support 8" Winchester/floppy or 5-1/4" Winchester and 5-1/4" or 8" floppy disk drives.



SMS

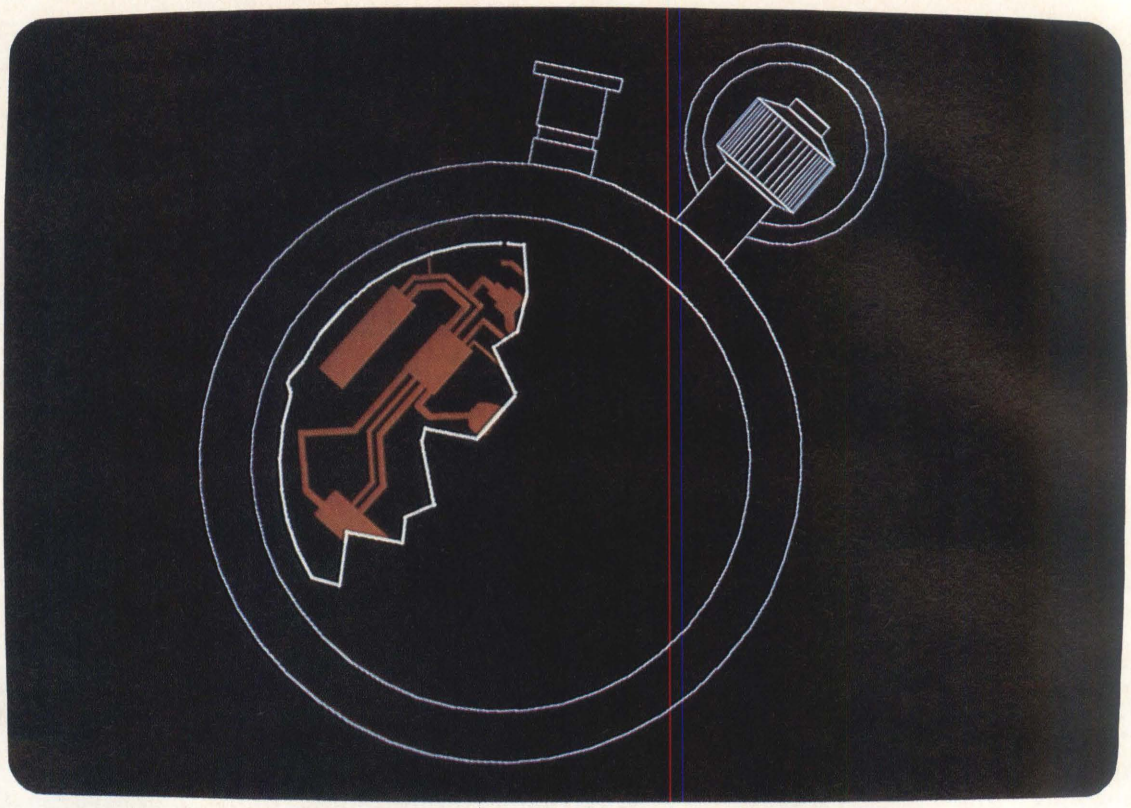
Scientific Micro Systems, Inc.

777 East Middlefield Road Mountain View, CA 94043
(415) 964-5700 TWX: 910-379-6577

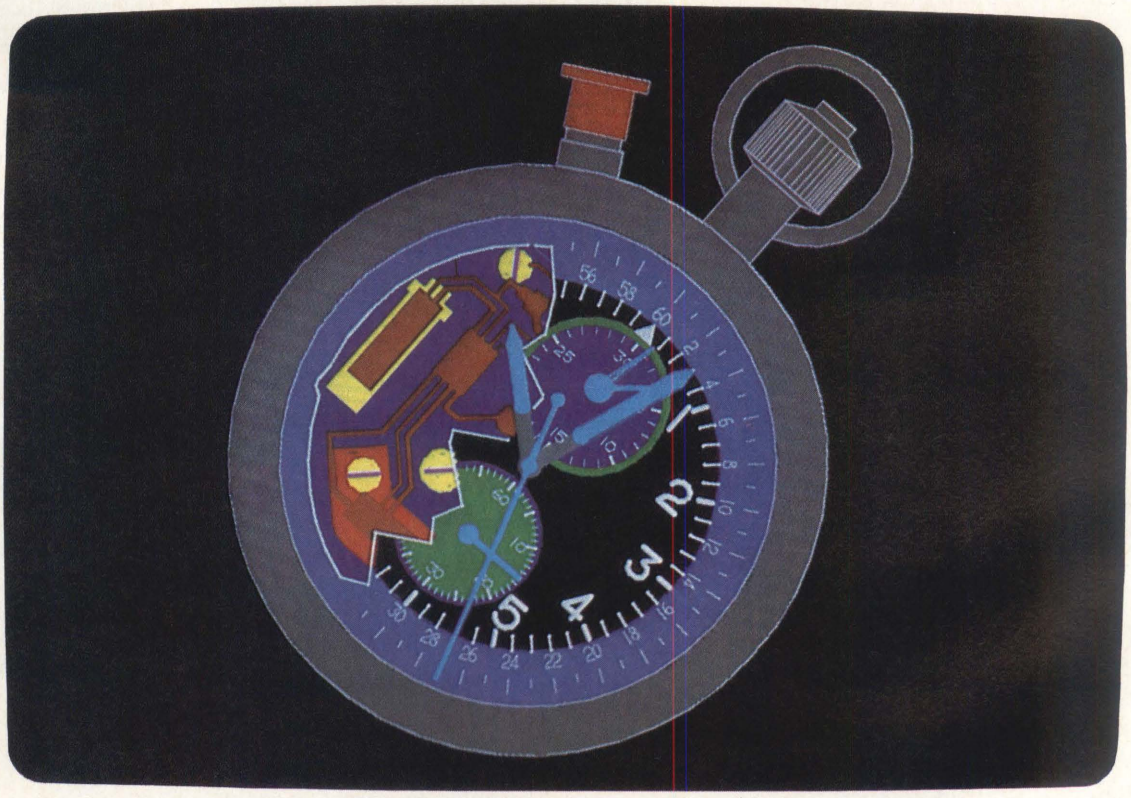
CIRCLE 109

SMS SALES OFFICES:
Seattle, Washington (206) 883-8303
Boston, MA (617) 246-2540; **Atlanta, Georgia** (404) 296-2029; **Morton Grove, Illinois** (312) 966-2711; **Yorba Linda, California** (714) 993-3768.

®TRADEMARK OF INTEL CORPORATION.



An average terminal gets this far in 1.4 seconds. You'll twiddle your thumbs for 14 seconds before it's finished.



New HiSCAN terminals complete the entire drawing in just over one second.

Graphics terminals with ten times the drawing speed. $2,195 Monochrome, $2,995 Color.

Stop watch courtesy of Heuer Electronic Corp.

"One, one thousand... done!"

Our new graphics co-processor technology gives you a faster draw than anything under $10,000. It's ten times as fast as others in its price range. A hundred times as fast as the slowest terminals.

You get superior resolution, too: 800 x 600 monochrome, 800 x 300 color. The monochrome terminal displays four gray levels. The color terminal displays 16 colors (out of 64 choices). Plus fifteen programmable, non-volatile function keys. Simple menus. Superior ergonomics.

You don't sacrifice a thing for superior graphics. The alphanumeric quality equals the best text terminals. The display is not interlaced (not running at half speed), so there's no smearing or ghosting when you scroll. You can even choose an 80- and 132-column display to get a full spreadsheet on the screen.

HiSCAN™ graphics terminals have full DEC VT220 and TEK® 4010/4014 compatibility. Plus your choice of DEC ReGIS, TEK 4027 or TEK 4105 compatibility at no extra cost. They're cable-ready for light pens, mice, digitizing tablets and inkjet printers.

They're designed to help you be more productive and creative. Because they work almost as fast as you think. So now you can play "what-if" with graphics.

You have to get your hands on a HiSCAN graphics terminal to believe it. And that's easy. We'll even arrange a 30-day trial at no risk. Just send the coupon.



New HiSCAN Terminals
From the people who brought you Retro-Graphics®

CD8/84

☐ Send complete information

☐ I want a 30-day trial with no risk. Call me about my compatibility requirements.

NAME _____ TITLE _____

COMPANY _____ PHONE _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____

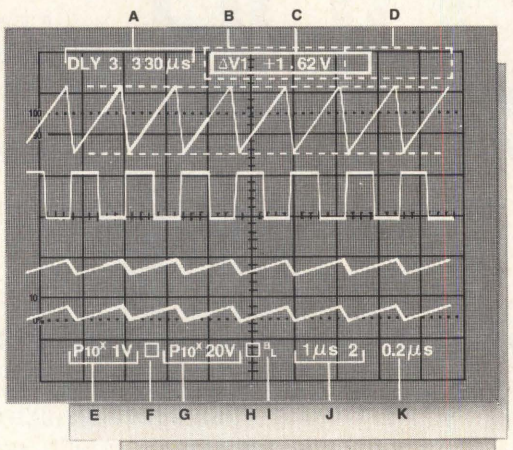
Mail to: Digital Engineering, 630 Bercut Drive, Sacramento, CA 95814

™ **DIGITAL ENGINEERING**
630 Bercut Drive, Sacramento, CA 95814
Phone (916) 447-7600 Telex 910-367-2009

© 1984 Digital Engineering, Inc. Retro-Graphics is a registered trademark and HiSCAN is a trademark of Digital Engineering, Inc.

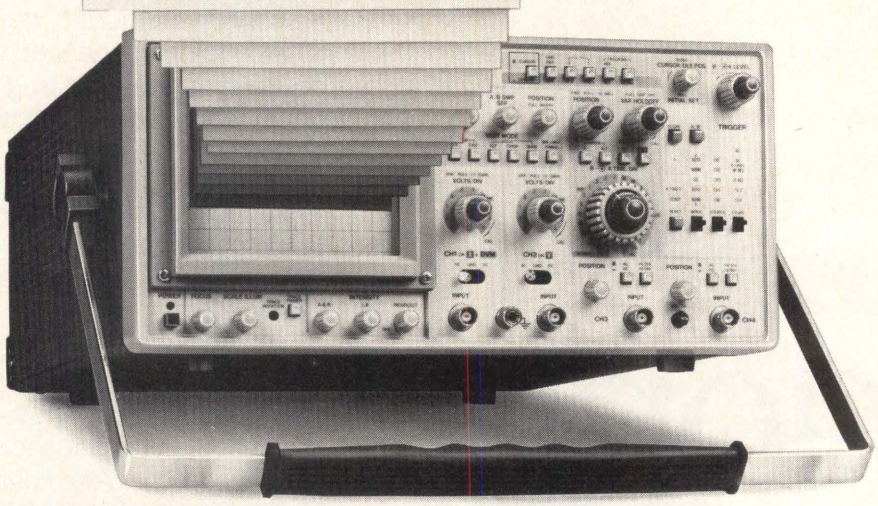
CIRCLE 110

THE INTELLIGENT SCOPE!



A. Delay Time
B. DVM Measured Values
C. Cursor Measured Values
D. DVM Relative Values And Frequency Measured Values
E. Scaling Factor (CH1)
F. Add Indicator
G. Scaling Factor (CH2)
H. Invert Indicator
I. BWL Indicator
J. Scaling Factor (Sweep A)
K. Scaling Factor (Sweep B)

HITACHI V-1100



The only Portable 100 MHz Oscilloscope with CRT readout and built-in micro-computer to make measurements faster, easier and more accurately.

Research, electronics, medical equipment development, military applications, or just general trouble-shooting...the Hitachi V-1100 Oscilloscope offers the ultimate in high-tech, high-versatility and high-value.

The Intelligent Scope is the smartest choice you can make.

This innovative state-of-the-art system is ruggedly engineered for enhanced reliability in the lab or out in the field.

Compare and you'll see why the Hitachi V-1100 is an incomparable value.

The exclusive Hitachi V-1100 features:

Readout functions	Digital Measurement functions	DVM	DC voltage: Corresponds to screen AC voltage: 50Hz to 10MHz	CH1 only	Display: 3 digits + units
		Frequency counter	Frequency: 120Hz to 99.9MHz		
	Cursor Readout function	REF, Δ cursors	Voltage Vabs: REF to GND Voltage V: Δ to REF Amplitude ratio: % (5 div = 100%) Time T: Δ to REF Time duty ratio: % (5 div = 100%) Phase: Deg (5 div = 360°) Frequency: 1/T = Δ to REF		
	Panel setting Value displays		Vertical axis: V/div, INVERT, ADD, BWL, UNCAL, MAG Sweep speed: s/div, UNCAL, MAG Other: delay time and A trigger source		
GND Reference function	Ground line display (GND LINE)		CH1 and CH2 displays are available.		

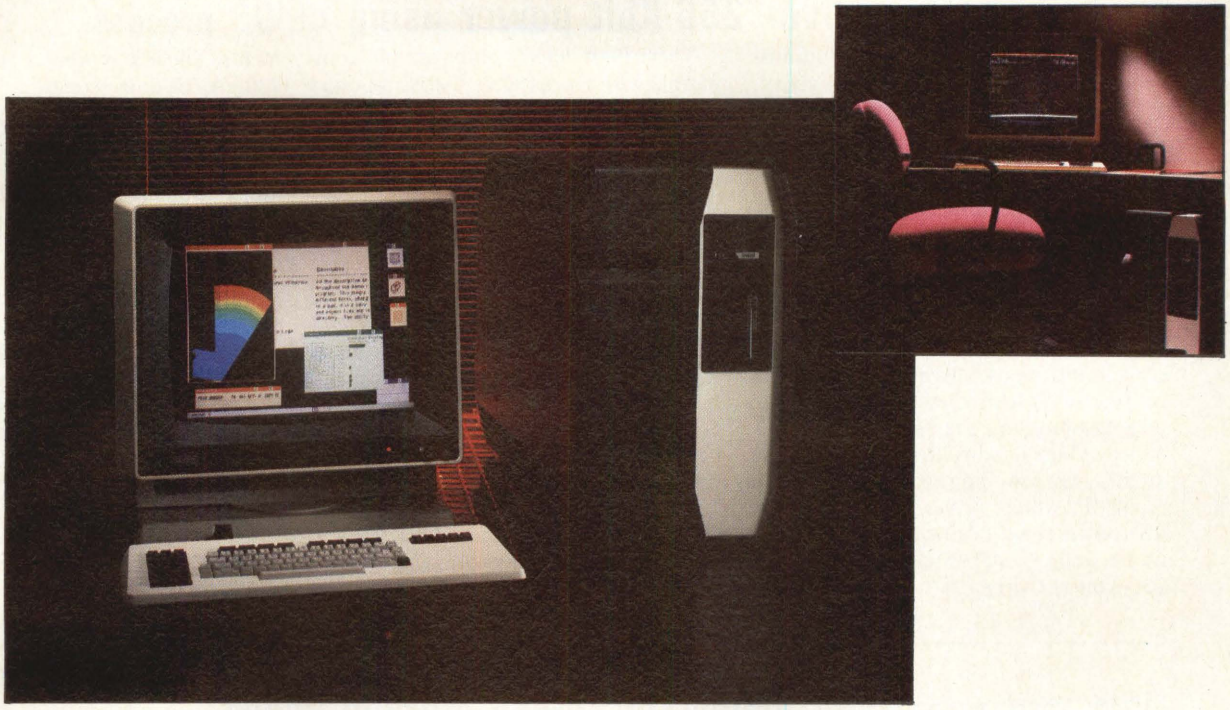
Additional features include quad channel (CH. 1, 2, 3, 4) with independent position controls, 8 trace with alternate sweep, 18kV-6" rectangular CRT, minimum deflection factor 1mv/div, maximum sweep time 2ns/div, TV-sync, X-Y operation up to 1 MHz (3° or less), variable hold-off, gate output for A and B sweep, CH 1 signal output to 100 MHz (-3dB)...plus much more.

HITACHI®

Hitachi Denshi America, Ltd.
175 Crossways Park West, Woodbury, New York 11797 ■ Telephone: 516-921-7200
For fast action, call TOLL FREE 800-645-7510

Chicago: 1725 North 33rd Avenue, Melrose Park, Illinois 60160 ■ 312-344-4020
Los Angeles: 18005 South Adria Maru Lane, Carson, California 90746 ■ 213-538-4880 (TOLL FREE) 800-824-9751
Dallas: 14169 Proton Road, Dallas, Texas ■ 214-233-7623

CIRCLE 111



Workstation graphics architecture lightens programming load

Compatible with Apollo's entire family of workstations and software, the DN550 color graphics workstation is 68010 based. The unit features a dedicated bit-slice graphics processor, a 1024 x 800 resolution color monitor, and a low profile detachable keyboard. Touch pad or mouse are available as options. The system offers up to 3 Mbytes of main memory and up to 2 Mbytes of double-buffered display memory that boasts 256 simultaneous colors.

The Graphics Metafile Resource (GMR) is a graphics architecture for the entire Domain workstation family. This software approach combines graphics capabilities with high graphics throughput and accommodates emerging standards such as the Graphical Kernel System (GKS), the Programmer's Hierarchical Interactive Graphics Standard (PHIGS), and the Virtual Device Metafile (VDM). GMR integrates a graphics data base with advance display routines in one package.

At the heart of the GMR architecture is the graphics metafile—a tree-structured data base that can be edited. This architecture stores graphic entities in the metafile, allowing information to be shared among applications and viewed on any workstation in the LAN. Built for speed, the virtual file stores up to 256 Mbytes of data directly in world coordinates, while the package uses its graphics hardware to increase throughput whenever possible.

Data in the metafile is displayed in multiple viewports within a given window, with the GMR handling scaling, translating, windowing, and clipping. Any changes are reflected in all viewports at the same time. Attributes can be associated with viewports or instances to control a range of display parameters including color, line style, and fill pattern. Hardcopy generation is simplified via the creation of a printable bit-map of a metafile.

The segmented data base supports nesting and instantiation. Nesting allows segments to contain other segments. Instancing reduces storage requirements for repetitive data. In addition, segments can store nongraphics data or reference data in other files.

Designed with a full 32-bit internal architecture, the unit's dedicated bit-slice graphics processor, together with extensive microcode, provides fast 2-D operations and more than a 10,000- transformed clipped vector/s throughput. In addition, the unit is capable of greater than 1-million pixel/s vector draws and 25- to 35- million pixel/s area fill and bit-block transfers. The system provides circle, arc, and spline generation with user-definable area fills and vector patterns.

With flicker-free color graphics and a multiwindow, multitasking environment, the unit allows up to 24 concurrent processes with 16 Mbytes of virtual address space/process as well as an integral interface to the 12-Mbit/s Domain LAN. Additional communication support comes in the form of IBM 3270 emulation, HASP, X.25, and an Ethernet gateway. Ergonomic features of the 19-in. display include both tilt and swivel and a nonglare filter.

Other available software includes D3M, which is a distributed database management system, SIGGRAPH Core graphics package, Fortran 77, ISO Pascal, C, and Lisp programming languages, as well as packages for professional support. Another package for the software engineering environment increases productivity in development, maintenance, and administration of software projects. In addition, terminal emulators, a font editor, and a high level debugger are available.

Options include a 50-Mbyte, 5 1/4-in. Winchester, a 45-Mbyte, 1/4-in. cartridge tape unit, four-slot Multibus peripheral adapter, and a floating point hardware accelerator. Configured with the Aegis operating system, the DN550, with 1 Mbyte of main memory, 4 planes of color, monitor, keyboard, and Domain LAN interface is priced at $31,500; $40,000 with a 50-Mbyte Winchester. **Apollo Computer Inc,** 330 Billerica Rd, Chelmsford, MA 01824. —*M.B.*

Circle 260

Dynamic RAMs deliver 256-Kbit power using CMOS process

The next generation design cycle loomsnear as 256-Kbit dynamic RAMs arrive. Billed as the industry's first CHMOS 256-Kbit DRAMs, the chips represent a memory increase that could redraw the boundaries of graphics, portables, and other applications.

Using its proprietary CHMOS technology, Intel employs dry etching and wafer stepping techniques to achieve NMOS performance with CMOS low power. These three 256-Kbit families appear on the scene just three months after introduction of Intel's 64-Kbit CHMOS DRAM. Basic chip versions are the 51C256H, which is aimed at high bandwidth applications; the 51C256L, a low power chip; and the 51C256HL, which strives to combine the low power/high performance attributes of its siblings. All three versions use a 256-Kbit x 1 organization.

These high density chips are pin compatible with the 64-Kbit HMOS 2164A and with the recent CHMOS 64-Kbit DRAM ranks. Thus, a fourfold increase in memory is accomplished without major changes in software or system design.

Effective bandwidth on the 51C256H and 51C256HL is enhanced by fast access Ripplemode operation. In one access, Ripplemode allows random read/write of up to 512 bits within a single row (see article, "CMOS 256-Kbit RAMs Are Fast and Use Less Power," p 133). This contrasts with the 4-bit "nibble" operation. Graphics displays are pegged as an immediate application requiring the bandwidth offered by the 51C256H and 51C256HL. The 51C256H has a 120-ns access time and a 65-ns Ripplemode cycle time.

The choice of CHMOS over NMOS offers the benefits of reduced power requirements. The low power 51C256L needs only 230 μA for data retention. This is about 1/20 the current required by comparable NMOS devices. CHMOS allows use of internal static circuits for improved addressability, as in the Ripplemode, without the power penalty that NMOS would incur.

These devices are currently available. Another 256-Kbit CHMOS DRAM—the 51C259, which features a static column address mode with 64-Kbit x 4 organization—is now available for sampling. Quantity-100 prices start at $115.45 for the 51C256H, $141.10 for the 51C256L, and $160.30 for the 51C256HL. **Intel Corp,** 3065 Bowers Ave, Santa Clara, CA 95051. —*J. V.*

Circle 261

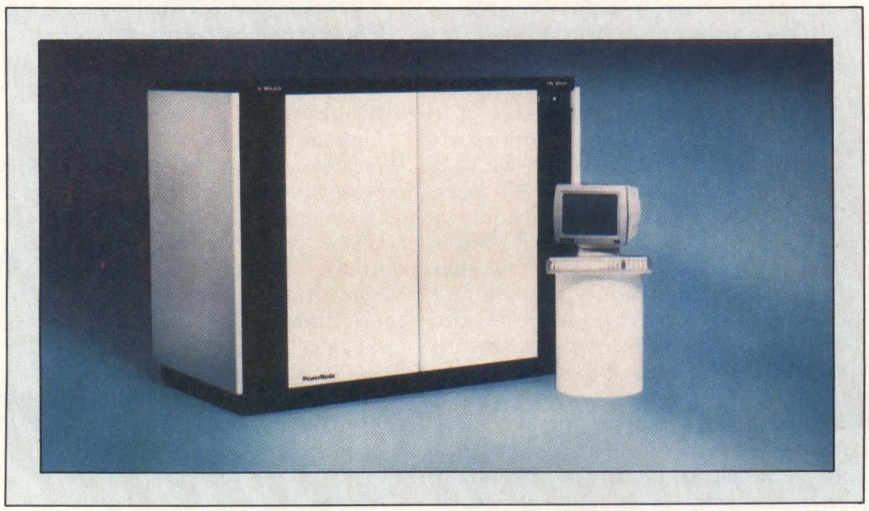
High end supermini melds Unix with marked speed

A 32-bit virtual memory processor capable of high speed performance combines with Unix to produce the PowerNode 9000 series. Benchmark ratings for the machines vault from 4.6 million instructions/s (MIPS) on the entry level PN9005 to 10 MIPS for the top of the line, twin processor PN9080. These performance levels are attained by using 10-K ECL technology, large cache memory, and four-stage instruction pipelining.

Optimized for Unix, the PN9000 CPU performs I/O, interrupts, and computational tasks. Based on 10-K ECL devices, the unit's four-stage pipeline allows completion of register-to-register instructions within one 75-ns cycle. In the high end 9080 model, a 32-bit internal processing unit (IPU) is added. The IPU boosts performance about 80 percent by handling computation-intensive tasks for the CPU.

Main memory directly links to the synchronous SelBUS, which transfers data at up to 26.67 Mbits/s. Series PN9000 units contain a minimum of four 1-Mbyte integrated memory module boards, built from high speed dynamic MOS RAMS with error correction and refresh logic. Memory interleaving—a built-in hardware feature—can efficiently distribute sequential addresses into independently operating memory modules.

The cache in the basic PN9000 processor includes 32 Kbytes of high speed memory. An optional cache memory module expands this capacity to 64 Kbytes. Since both the IPU and CPU have their own cache memories, the cache for a fully configured PN9080 can reach 128 Kbytes. The PN9080 supports up to 16 Mbytes of main memory.



Other PowerNode 9000 features include an integral floating point accelerator, instruction and operand prefetch, and base register support. The units support up to 16 Mbytes of virtual memory. Optional multiply accelerators speed execution of single- and double-precision floating point multiply operations.

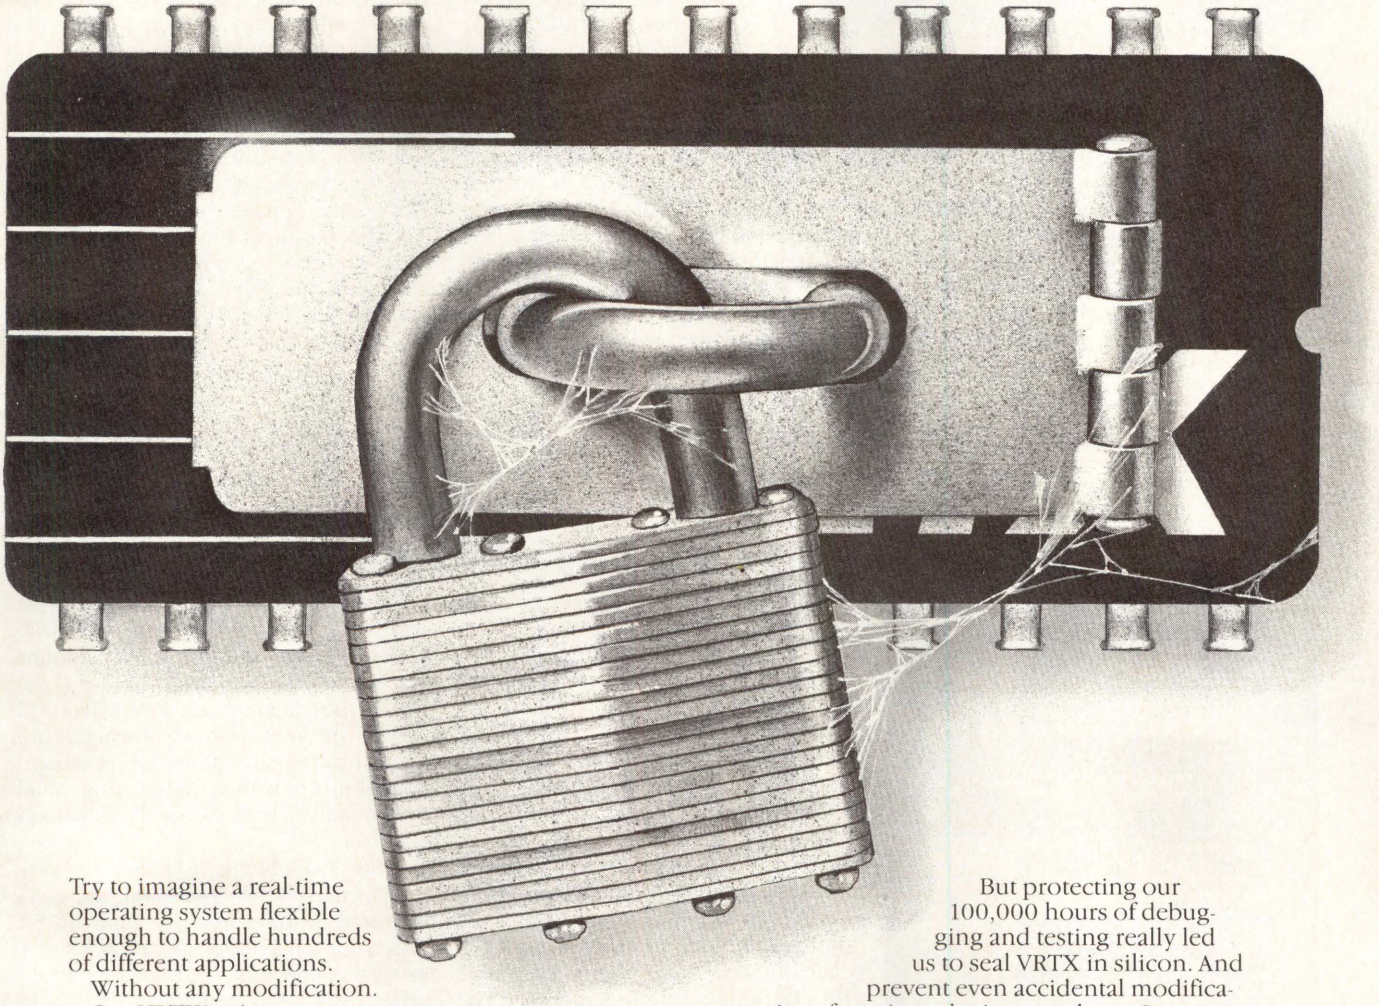
The machines feature the UTX/32 operating system—Gould's authorized version of Unix. This system is based on BSD 4.1C, with selected features from AT&T's Unix System V. A line of Compatibility Suite application software, including Unify Corp's Relational Database Management System, is available.

PowerNode units support multiprogramming and multiple users through local or remote terminals. These can stand alone or can be used in a LAN. The Transmission Control Protocol/Internet Protocol (TCP/IP) standard, RS-232 and high speed parallel port links, and the Ethernet 1.0 specification are supported.

The entry level PN9005, which includes CPU, 32 Kbytes of cache memory, and 4 Mbytes of main memory, costs $245,000. The PN9080 is priced at $385,000. **Gould Inc, Computer Systems Div,** 6901 W Sunrise Blvd, PO Box 9148, Fort Lauderdale, FL 33310. —*J. V.*

Circle 262

All this success and we haven't changed a bit.



Try to imagine a real-time operating system flexible enough to handle hundreds of different applications.

Without any modification.

Our VRTX® microprocessor operating system can do just that. Over 200 diverse applications, from navigating aircraft to controlling disks, to playing video games, have been successfully implemented around VRTX. All without tweaking a single bit of VRTX code.

A good thing, too.

Because our VRTX operating system is delivered in 4K bytes of ROM. Which naturally makes tweaking a little impractical.

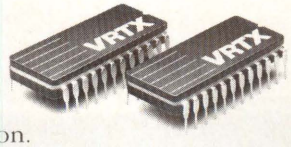
But it does make VRTX the most bug-free operating system you'll ever use. And the easiest one you'll ever install.

In fact, we wouldn't be surprised if you saved six to 12 months of development time using VRTX.

But protecting our 100,000 hours of debugging and testing really led us to seal VRTX in silicon. And prevent even accidental modifications from introducing new bugs. So now we know, no matter how successful we get, we'll never get spoiled.

For a free VRTX evaluation package (including timings for system calls and interrupts) contact us with the details of your application, including the microprocessors you're using: Z8000, Z80, MC68000 or 8086 family. Write Hunter & Ready, Inc., 445 Sherman Avenue, Palo Alto, California 94306. Or call (415) 326-2950.

HUNTER ◆ READY

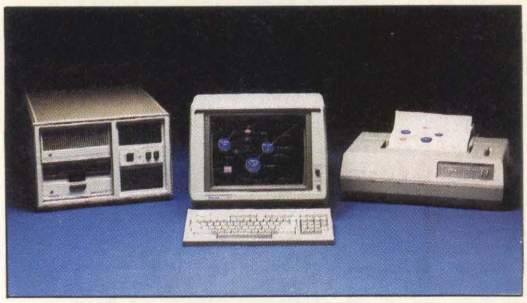


Operating systems in silicon.

© 1983 Hunter & Ready, Inc.

CIRCLE 112

Structured analysis eases project specification



Techniques of structured analysis developed in data processing are used in a set of software tools tailored to meet the needs of software project specification. The tools are designed to define system and software problems at the beginning of the design cycle, before the actual code is written. These Structured Analysis (SA) tools cover four main areas: graphics diagram editing, internal consistency checking, error correcting, and formatting the analysis for output on printers or graphics peripherals.

To achieve system specification, a color graphics terminal displays data flow paths, processes that modify data, and data storage elements. Data flow diagrams are arranged in the computer hierarchically, enabling the user to access each branch of the tree-like structure. At the bottom of the hierarchy is the mini-specification—a brief, English-like description of a data-modification process. Also included is a data dictionary that defines the items in the flow diagram and in the mini-specifications. The main function of the system is to track activity and notify the designer of any loose threads, such as dead-end data paths or undefined processes.

The SA graphics editing tools run on a variety of Tektronix color graphics terminals. The system's internal consistency tools keep track of the processes and data paths that the user enters. These routines automatically check for errors in data flow diagrams, mini-specifications, and the data dictionary. They also search out inaccuracies in syntax, conflicts in parent-child relationships occurring within the data flow hierarchy, and other problems such as repeated data names.

Currently, SA tools run on the Tektronix 8560 microcomputer development system and on the DEC VAX. Versions for VAX-Unix environments and for VAX-VMS are also under development. Hardcopy output compatible with Tektronix printers and plotters, and with the 4695 color copier, is available. **Tektronix, Inc,** PO Box 1700, Beaverton, OR 97077. —*T. W.*

Circle 263

Unix systems provide long-term solutions with upgradeability



A family of 68000-based systems for 1 to 16 users, the very intelligent Unix system (VIX), offers a 9-slot Multibus card cage (optionally upgradable to 16 slots). Model I has two 5 1/4-in. dual-sided, dual-density disk drives with 1.4-Mbyte capacity combined with 5 1/4-in., 20-Mbyte Winchester drive. Model II has 640-Mbyte disk storage capacity and 20-Mbyte streaming tape drive. Supported languages include C, Fortran 77, Pascal, Basic Plus, SMC Basic, Macro assembler, and simple assembler.

Both models feature UniPlus, a Unix System III operating system derivative with Berkeley enhancements. The enhancements provide a general purpose Unix system for standalone computing, distributed data processing, and communication. Written in C, the operating system includes a C compiler and simple assembler.

Using the 8-MHz 16/32 68000 micro with memory protection, the systems have 256 Kbytes of RAM on board (expandable to 1024 Kbytes), interrupt select logic, and an MMU. The Multibus card cage holds the CPU/main memory, main peripheral, RAM, I/O system boards, and mass storage modules.

As a companion to the CPU/memory board, the main peripheral module contains two serial programmable RS-232 interfaces, up to 64 Kbytes of EPROM, and an internal timer. Processor options are accessible via DIP switches.

The RAM memory module offers 512 Kbytes of RAM with error correction, while the mass storage module is a multifunction controller supporting dual floppy and Winchester disk drives and streamer.

Also available is a series of very intelligent systems combining universal terminal emulation, 8- and 16-bit processing capabilities, as well as CP/M-80 or MS-DOS operating systems. The 8200 VIS series uses the company's intelligent terminals. Both terminals and systems support emulations for IBM 3275/71/77, DEC VT 100/200, and Data General's 410/460. Additional features include an 8- or 10-MHz 68000 with 1024 Kbytes of RAM, 128 Kbytes of EPROM, and 2 Kbytes of EAROM.

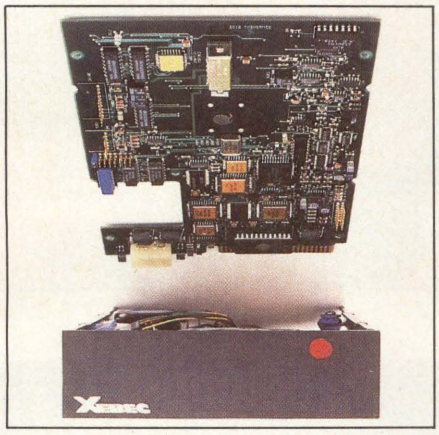
Prices range in moderate volume quantities from $10,000 for Model I to $23,000 for the Model II, not including terminals. A VIS, in a typical configuration, is $4350 in moderate volume quantities. **Megadata Corp,** 35 Orville Dr, Bohemia, NY 11716. —*M.B.*

Circle 264

Winchester merges with controller for single-board integration

The Owl 10-Mbyte Winchester disk drive incorporates a S1410A disk controller on the drive electronics board. Previously, one board held the controller and one had the drive. The intelligent half-high 5 1/4-in. device is suited for single-user micro-based systems, as a built-in mass storage device, or as an add-on storage subsystem.

Board components are a mix of both standard LSI and custom devices. All are surface mounted, thus making possible the combination of two traditional 5 1/4-x 8-in. boards (disk electronics and controller) on one board having the same dimensions. The actual board footprint is 1.63 x 5.75 x 8.0 in. Using custom circuits, all S1410A disk drive controller functions can be implemented with fewer devices.



One chip combines the serialization, deserialization, and error detection/correction functions with a digital data separator. The built-in data separator is optimized for increased read and write performance. Ground-loop noise and induced errors are reduced via the closeness of the data separator to the drive. The LSI components include DMA, a Z80, a programmed I/O chip for the Z80, and the SASI interface component.

Providing the high level interface between the disk drive and host, the controller interface is made up of a generic command set plus an 8-bit parallel data bus for communication. This high level command set reduces host system overhead and is compatible with all Xebec SASI systems. A system integrator needs only to design an SASI host adapter onto the host to accommodate the Owl drive.

The built in controller provides automatic seek and position verification and automatic command retry when an error occurs. Other features include multisector data transfer with automatic cylinder and head switching, and programmable sector interleave.

The oxide media drive's 10-Mbyte formatted storage capacity has a 99-ms average access time. It uses a rack and pinion actuator and open-loop stepper head positioning . Other features include head shipping zone, MFM encoding, and a 5-Mbit/s transfer rate.

To increase reliability further, the matched components share one power supply, lowering induced error rates. In addition, traditional disk drive cabling is eliminated, reducing transmission line effects.

In quantity, the drive/controller combination is under $500. **Xebec,** 432 Lakeside Dr, Sunnyvale, CA 94088.

—M.B.

Circle 265

Gate array series reaches 11,000-circuit density

Using CMOS technology, the μPD65100 achieves 11,000 gates—surpassing the previous industry high of 8000 gates. The semicustom ICs ensure flexibility by means of a manufacturing process with two-level metallization. With 90 percent gate utilization and maximum power dissipation of 20 μW/gate, the arrays are packaged in 72, 132, 176, or 208 pin grid arrays.

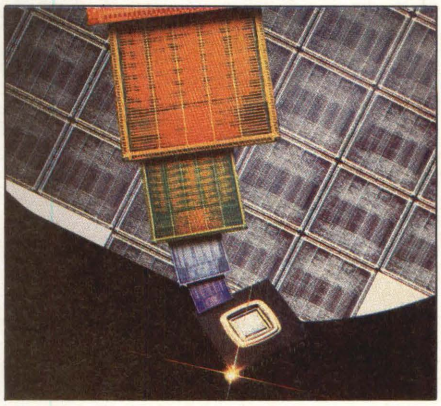
Other gate arrays in the series include an 8000-gate version, a 3300-gate version, and a 2100-gate version. All are manufactured with two millimicron channel lengths. Processing speed is 2 ns/gate with three fan outs using 3-mm wiring.

The gate array development process is structured so that users will not be restricted to one interface type. The simplest interface requires a circuit diagram and test patterns. The company then takes on the rest of the development process. Other available interfaces range from providing a net list and test pattern in NEC-compatible format to graphics PC-generated data, where information is generated via the PC9800 workstation. This workstation supports schematic capture and limited design rule checking.

In this gate array series, the entire development process begins with design rule checking. Parameters such as cell usage, power dissipation, and fan-out loading are determined and checked. Then, all coding errors and data conversion errors are checked and eliminated. Prior to automatic placement and routing, a delay time simulation provides an expected circuit delay analysis. At this point, placement and routing software allows up to 95 percent cell utilization without resorting to manual routing.

Further in the development process, a final delay time simulation occurs. Actual wire lengths are taken into account for accurate circuit analysis. If all previous steps are successful, the design enters production, followed by testing of all wafers. Wafers are divided into individual chips that are die bonded onto customer-specified packages. Each chip is wire bonded and sealed with dc parameters and its logic functions are checked by a final test.

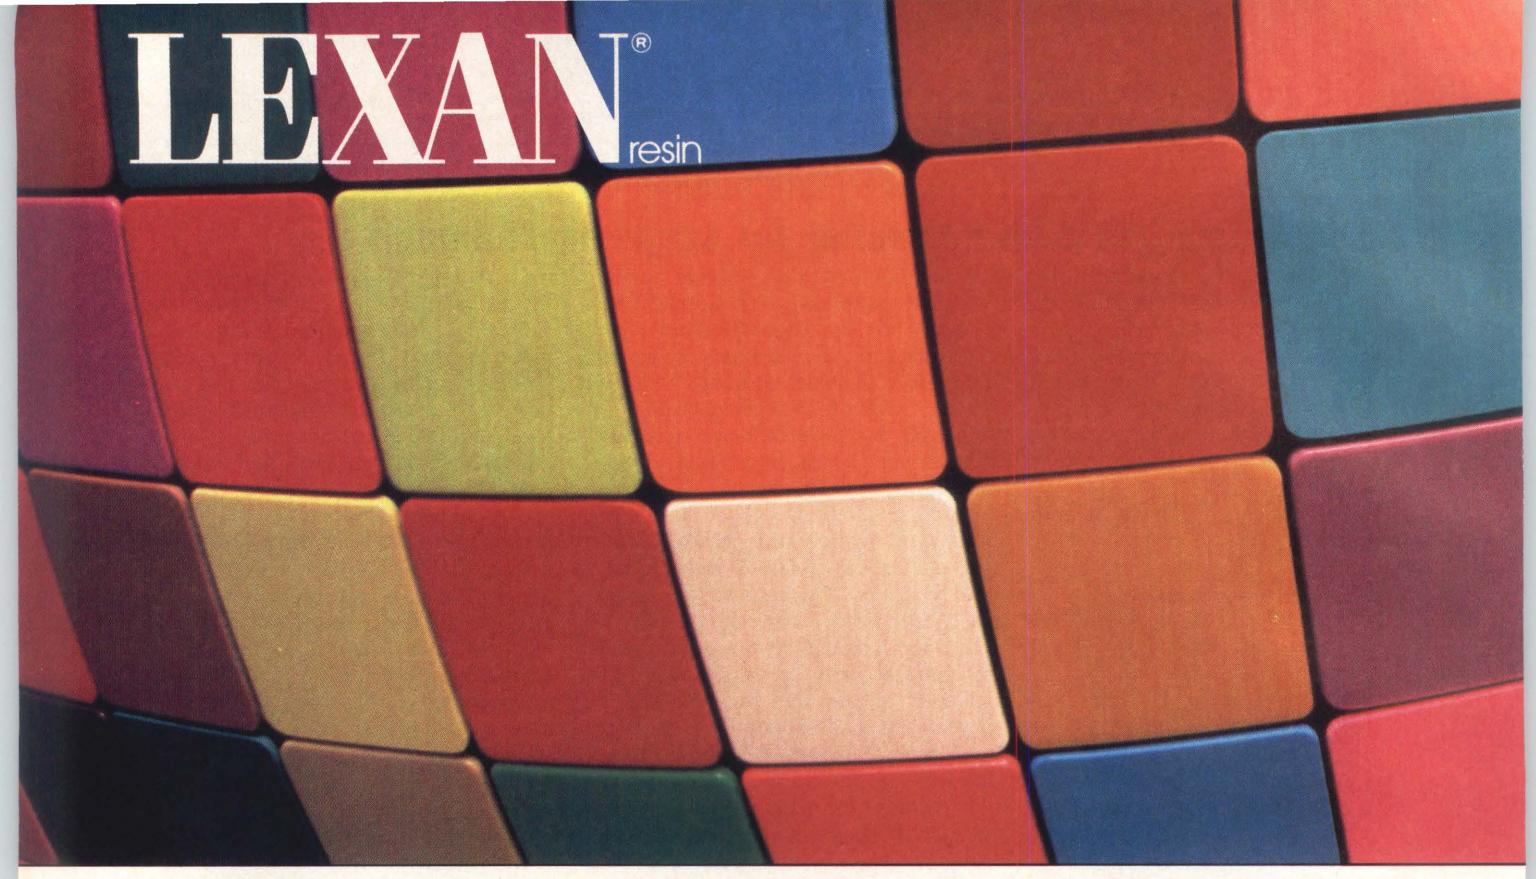
In addition to the CMOS arrays, an ECL type gate array with 3000 gates, a delay time of less than 1 ns/gate and power dissipation of 1.1 mW/g is available. Two



other ECL gate arrays with subnanosecond delay times, and one TTL with a 2-ns/gate delay time round out the gate array offerings.

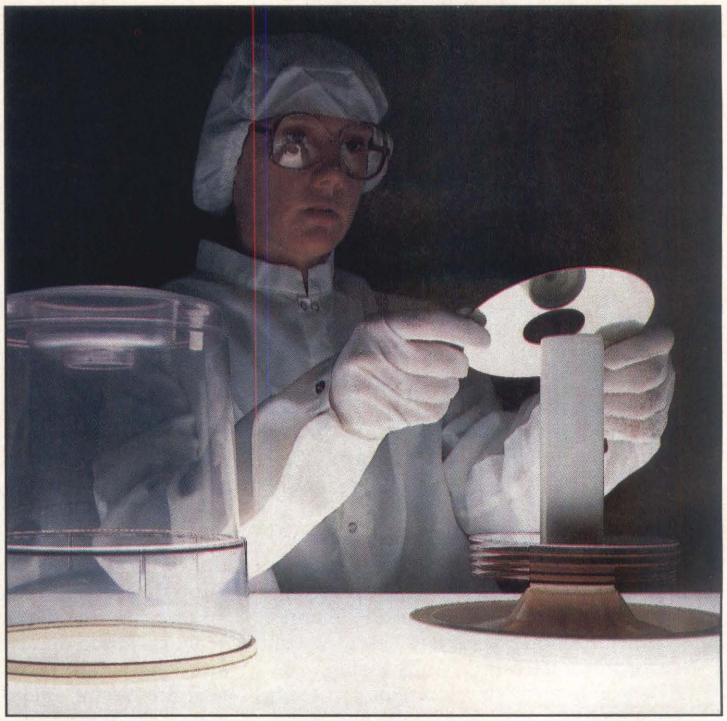
Pricing varies according to package type, but an approximate figure is $200 per unit in 5000-piece quantities plus an $80,000 nonrecurring engineering option. Delivery is approximately 10 to 12 weeks after initial simulation and verification. **NEC Electronics Inc,** 401 Ellis St, Mountain View, CA 94043. *—M.B.*

Circle 266



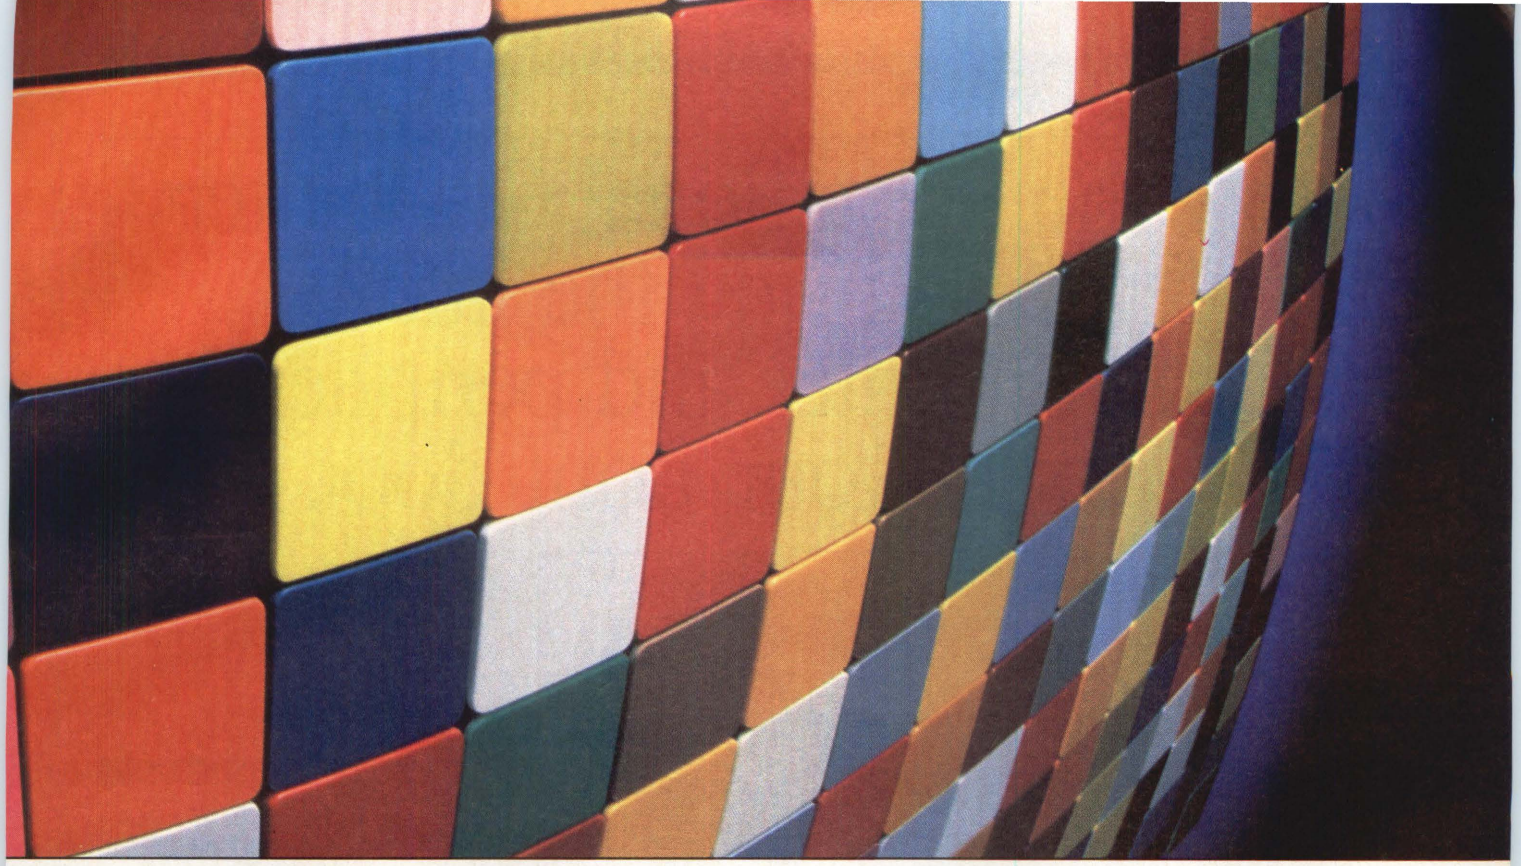
It's 19,946* colors.

It's crystal clear.



LEXAN® resin

* As of June 1, 1984.
® Registered Trademark of General Electric Company.



It's whatever you need it to be.

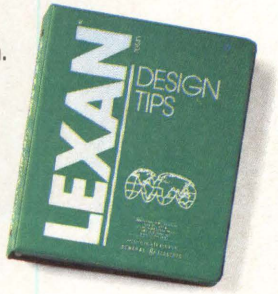
Custom Colors. Available in over 6,000 active production colors, the world's premier engineering plastic is every bit as eye-catching as it is high performing. In fact, we've already matched nearly 20,000 colors in the LEXAN resin family—colors that are functional or decorative, opaque, translucent or transparent—and always consistent.

Clarity Plus. LEXAN resin also offers clarity, dimensional stability and impact strength competitive materials can't match. That's why it was chosen for the first reusable, all-plastic disc storage container. The result is a portable "clean room" for computer disc protection during production and shipping.

Constant Quality and Quality Control. All colorants are screened for stability, consistency, aging properties and compatibility with LEXAN resin. And since there's always more to learn about the color process, our color specialists are members of the National Bureau of Standards' Manufacturers Council on Color and Appearance.

Bottom Line Value. Explore LEXAN resin's design opportunities and processing advantages, right from the initial design, to lower your finished part cost.

Brighten up your next application with LEXAN resin. For immediate literature, and to receive a free subscription to the new LEXAN **Design Tips,** simply dial (800) 422-1600; in Vermont, (802) 442-8356.



General Electric Company
Plastics Group
LEXAN Products Division
One Plastics Avenue
Pittsfield, MA 01201

We bring good things to life.

GENERAL ELECTRIC

CIRCLE 113

OEM Decision Makers

"Only the Invitational Computer Conferences bring the latest OEM computer and peripheral products to your front door. You'll find us there!"



BEN WANG, President
Wangtek

And you'll find other top OEM manufacturers, such as IBM, Control Data, DEC, Fujitsu, NEC and Seagate, to name a few.

In their 14th year, the "OEM Only" Invitational Computer Conferences bring you, the volume buying decision makers, together with the key suppliers of computer and peripheral products. The ICCs, a series of ten, one-day regional shows are convenient to where you live and work. The social business setting makes it easy for you to meet potential suppliers one-on-one, and attend high tech seminars of your choice. As an invited guest, there is no cost to you.

Hear what the OEM manufacturers have to say, learn more about their products, and remember, you may attend "by invitation only."

1984/85 U.S. ICC Locations

Sept. 6, '84	Newton/Boston, MA
Sept. 25, '84	Southfield/Detroit, MI
Oct. 10, '84	Cherry Hill, NJ
Oct. 23, '84	Englewood/Denver, CO
Jan. 8, '85	Irvine, CA
Jan. 29, '85	Houston, TX
Jan. 31, '85	Dallas, TX
Feb. 26, '85	Ft. Lauderdale, FL
Mar. 19, '85	Palo Alto, CA
Apr. 2, '85	Nashua, NH / No. MA

Call your local OEM supplier for your invitation or fill out the coupon and mail to:

B.J. Johnson & Associates, Inc.
3151 Airway Ave., #C-2
Costa Mesa, CA 92626
Phone: (714) 957-0171
Telex: 188747 TAB IRIN

Yes! I need an invitation to your "OEM Only" ICC. The nearest ICC to me is: _____

I buy in volume:

- ☐ Computers
- ☐ Disk/Tape Drives
- ☐ Controllers / Interfaces
- ☐ Terminals / Graphic Displays
- ☐ Software
- ☐ Printers
- ☐ Memory Boards
- ☐ Modems / Multiplexers
- ☐ Power Supplies

Name _____

Title _____

Company/Division _____

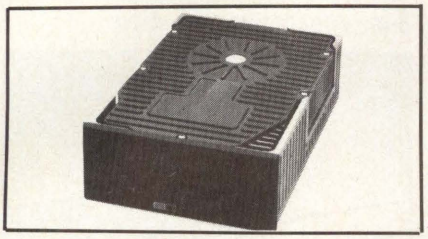
Address _____

City _____ State _____ Zip _____

Mail To: B.J. Johnson & Associates, Inc., 3151 Airway Avenue, #C-2, Costa Mesa, CA 92626
Phone: (714) 957-0171 Telex: 188747 TAB IRIN

CIRCLE 114

Rugged 3.5-in. Winchesters can withstand shock of 100 G



A four-point internal suspension system allows the Ranger family to survive a 40-G shock with power off and 100 G with the addition of external shock absorbers. A proprietary head lifter and arm lock fully protect the heads and media from damage. Model 3521 and 3522 offer formatted storage capacities of 5 Mbytes (one disk) and 10 Mbytes (two disks), respectively, in a half-height configuration. Equipped with the ST-412/506 interface, the drives have a 65-ms average access time. **LaPine Technology,** 1111 Space Park Dr, Santa Clara, CA 95050.
Circle 267

Low cost GCR magnetic tape system attaches to HP 1000 series

The Series Ten features fast dual-density 800/1600 chars/in. or tri-density 800/1600 at 75 in./s and 6250 chars/in. at 45 in./s. It requires one slot in the chassis and is plug compatible with the HP operating system and DVR23 driver. Transfer rates are up to 280 kbytes/s. Features include single-board interface on M/F/E chassis and up to four transports of the same speed. Dual-density version is priced at $17,390 and tri-density is $19,995. **Dylon Data Corp, div of Integritek, Inc,** 9561 Ridgehaven Ct, San Diego, CA 92123.
Circle 268

High density Q-bus memories sport 512-Kbyte and 2-Mbyte capacities

Designed for the LSI-11/73 and the MicroVAX, Q-RAM 11B and Q-RAM 44B, memory products feature an 80-ns read and a 45-ns write access time, as well as block mode DMA. The Q-RAM 11B uses 64-Kbit dynamic RAMs to achieve 512-Kbyte capacity, while the Q-RAM 44B uses 256-Kbit DRAMs to garner 2-Mbyte capacity. Onboard control and status register, as well as LED parity error indicators, are standard. **Clearpoint, Inc,** 106 South St, Hopkinton, MA 01748.
Circle 269

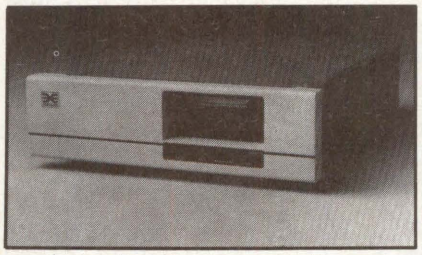
Thin-film recording heads dense-pack data

Cyber 300 thin-film recording heads, developed for 14-in. Winchester applications, contain an inductive read/write thin-film transducer. The devices operate at densities up to 800 tracks/in., 10,000 flux changes/in. on 3380-type oxide media. Cyber 300s are shipped as head gimble assemblies or head arm assemblies. Minimum order size for evaluation quantities is 10 heads for $1000. **Cybernex Corp,** 6580 Via Del Oro, San Jose, CA 95119.
Circle 270

Dynamic RAM module uses 24-bit addresses

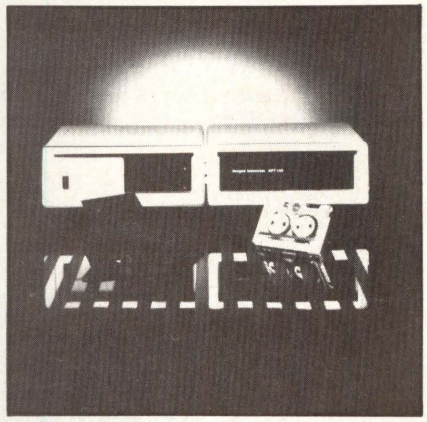
Designed to accept 72 64-Kbit dynamic RAMs, the MD-4539 module features 24-bit addresses and 8- or 16-bit data. It has a typical access time of 240 ns and offers 4-Kbyte units that allow free mapping within a 16-Mbyte range. Parity generation and checking are standard. A four element LED array displays error flags. For prohibitive addresses, the MD-4539 offers a maximum of two 42-Kbyte and two 64-Kbyte areas that may be write protected. **Symbicon Assoc, Inc,** 89 Rte 101A, Amherst, NH 03031.
Circle 271

Disk and tape subsystem enhances storage/backup



Medley subsystem combines a 5¼-in. Winchester disk with a CDC Sentinel quarter-inch cartridge tape drive, employing SCSI to provide enhanced storage/backup for DEC Q-bus and Unibus computers. Medley comes in either a desktop cabinet or standard RETMA rackmount. The Winchester offers 36- or 110-Mbyte formatted data storage, while the cartridge tape drive delivers up to 70 Mbytes of backup. Price for the 36-Mbyte version is $9795 for the Q-bus, and $10,195 for the Unibus. **Emulex Corp,** PO Box 6725, 3545 Harbor Blvd, Costa Mesa, CA 92626.
Circle 272

High capacity backup is MS-DOS, PC-DOS compatible



Super Performance 65- and 140-Mbyte hard disk drives have gained a high capacity backup option. Designated the Dragon 3.3-Mbyte Super Floppy Disk Drive Backup Option Model SPF, this device has a 500-kbit/s transfer rate and 3-ms track-to-track performance. MS-DOS and PC-DOS compatible, it can read 48- and 96- track/in. diskettes. It sells for $995. A 60-Mbyte, quarter-inch streamer tape backup is also available. **Dragon Industries,** 35 Main St, Hopkinton, MA 01748.

Circle 273

Streaming tape subsystem offers archival storage

An 8-in. streaming tape subsystem handles system file backup and archival storage for fixed Winchester and removable disk media. The subsystem provides 45 Mbytes of formatted data backup in less than 10 min on a single cartridge. It uses ANSI standard quarter-inch tape media and compiles with QIC-02/4 interface and format standards. The subsystem is priced at $3395. **Naked Mini Div, Computer Automation,** 18651 Von Karman, Irvine, CA 92713.
Circle 274

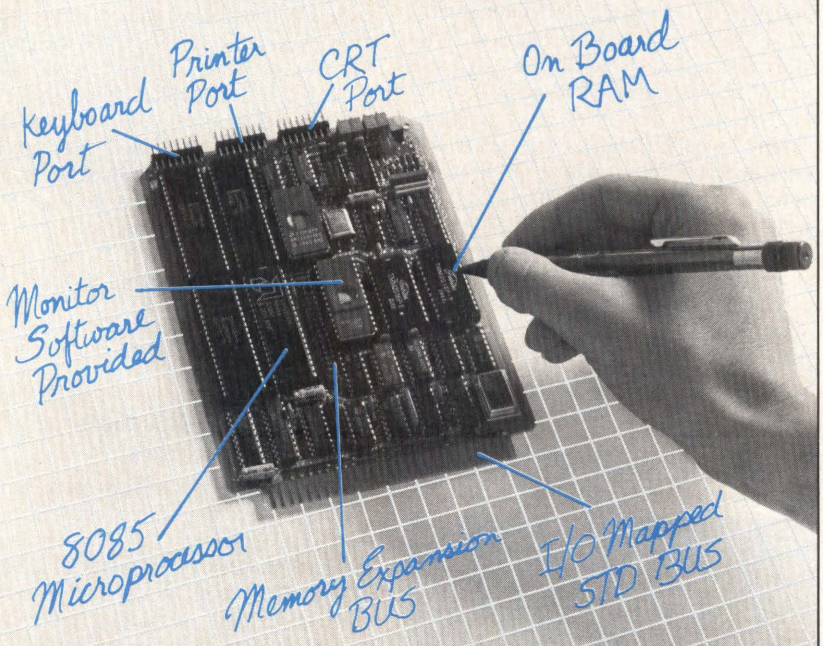
Full-height Winchesters conform to industry standard dimensions

The 5000 series 5¼-in. Winchesters have capacities that range from 6.38 to 25.5 Mbytes with one to four disk platters per drive. All drives are individually tested and inspected, including a 48-hour high temperature burn-in. MTBF is in excess of 15,000 hours. **Disc Tech One, Inc,** 849 Ward Dr, Santa Barbara, CA 93111.
Circle 275

Memory board mixes dual-height with block mode DMA

In sizes of 256 Kbytes, 512 Kbytes, and 1 Mbyte, the SMSV11 series supports the DEC block mode DMA bus protocol on systems with 18- or 22-bit addressing. It has a single 5-V low power requirement. Two such boards can sit side-by-side on a quad backplane. **Webster Electronics Inc,** 333 Cobalt Way, Suite 106, Sunnyvale, CA 94086. **Circle 276**

SMART CRT CONTROLLER ON STD BUS



Cubit's new I/O Processor controls a CRT, printer and keyboard.

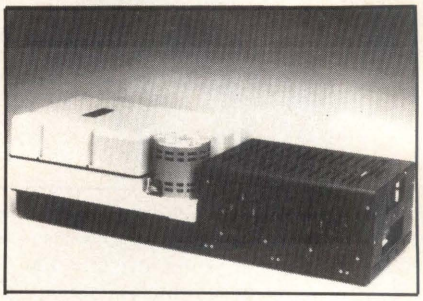
Using an 8085 microprocessor with its own memory, the board frees your system CPU to race ahead of slower peripherals. Terminal-like commands permit easy communication between this smart controller and the host-processor.

Bring distributed processing to your STD Bus system for $345 in single quantity. With stock to two week delivery, you won't have to wait long.

CuBIT INC.
Division of Proteus Industries

190 South Whisman Road, Mountain View, CA 94041
Telephone: (415) 962-8237

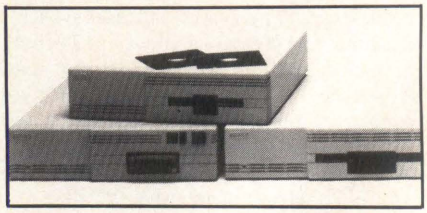
Large parallel RAM packed in 8-in. Winchester disk drive



An 8-in. parallel RAM Winchester disk drive holds 212 Mbytes of unformatted storage on eight data surfaces that divide into two sets of four parallel read/write channels. The MVP212 consists of two assemblies: a sealed head-disk assembly containing five disks, eight heads, and a rotary voice-coil positioner; and an electronic module containing power amplifier, servo tachometer, microprocessor control, I/O, and four read/write PC boards. Price in quantity is $7500. **MegaVault,** 6431 Independence Ave, Woodland Hills, CA 91367.
Circle 277

Disk controller firmware simplifies HP file sharing

Option 660 firmware enhances Series 3000 Winchester disk subsystems by allowing file sharing among two to three Hewlett-Packard computers. This firmware, residing in the controller, arbitrates use of a common data storage area in the hard disk. Thus, applications that do not require a full-scale LAN can gain connectability. Available sizes of shared information space are 1.2 Mbytes, 4.9 Mbytes, and 9.8 Mbytes. Connection is via standard HP-IB cables and, for distances less than 66 ft. (20 m), no additional hardware is needed. Quantity one price of Option 660 is $480, with a two-port option for $700, and a three-port option for $850. **Bering Industries, Inc,** 1400 Fulton Pl, Fremont, CA 94539.



Circle 278

Microbar. Your future system is our current project.

At Microbar, we're working today on the Multibus™ single-board computer (SBC) you'll need tomorrow. We specialize in developing SBCs that put OEMs and system integrators in front of the competitive pack—in technology, in time-to-market, and in cost. Doing that means maintaining a clear focus on the future and anticipating your needs—maybe even before you do.

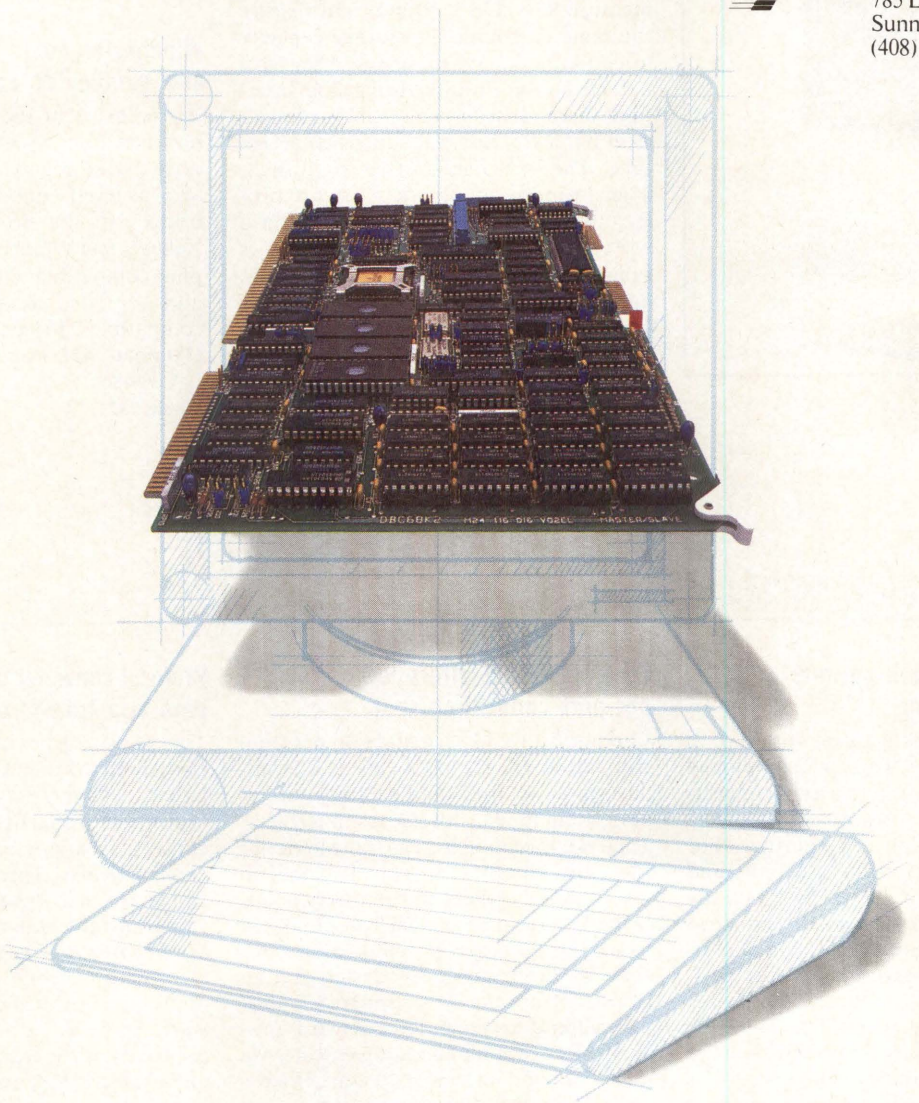
Keeping you ahead means we've got to stay flexible. We refuse to lock in to particular CPUs and technologies. Because we're independent, we continuously—and objectively—evaluate the newest microprocessors, busses and operating systems. So you can count on getting the best fully-integrated SBC solution for your application. And on getting it early enough to make that critical difference.

Today, we're exacting maximum performance from the leading 16- and 32-bit microprocessors implemented with the industry-standard Multibus. We do it by applying systems engineering concepts to board-level design. And through special architectures leading to speed and efficiency far beyond typical Multibus capabilities.

One thing we don't—and won't—do is compete with *you*. Our business is high-performance CPU boards, not systems. So our resources are 100% committed to developing and producing SBCs that make *your* systems successful, that give you a real edge in the market.

Microbar. We're working on your next system—today.

MICROBAR SYSTEMS, INC.
785 Lucerne Drive
Sunnyvale, CA 94086
(408) 720-9300

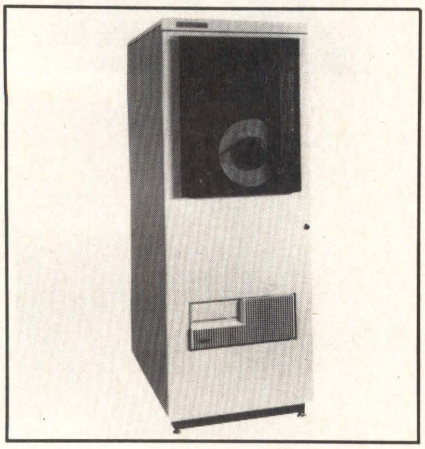


™Multibus is a trademark of Intel Corporation.

CIRCLE 116

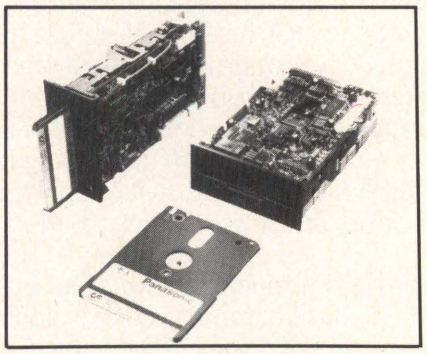
Disk and backup tape combine in small footprint unit

A 132-Mbyte Winchester disk drive pairs with the HP 7974 tape drive in a 5-ft (1.6-m) cabinet to become the HP 7914ST. This cabinet, which also holds a controller, has room for a computer as well. Options include software-selectable 800- and 1600-char/in. half-inch tape; quarter-inch streaming 67-Mbyte cartridge tape, and a dual controller that enables HP 3000 computers to communicate simultaneously with both the quarter-inch cartridge tape and the disk. The standard HP 7914ST costs $26,000. **Hewlett-Packard Co**, 1820 Embarcadero Rd, Palo Alto, CA 94303.



Circle 279

Compact 3-in. floppy drives and media store from 250 Kbytes to 1 Mbyte



There are six models in the family with designations EME-102/202, 150/250, and 130/230 and all are plug compatible with popular 5¼-in. interfaces. The drives measure 99 x 150 x 40 mm and feature the same or double the storage capacity of conventional 5¼-in. drives. They have a single button for diskette ejection/insertion, direct drive brushless motor, and a steel band for fast track-to-track access time. The 3-in. media offers rigid shell, a head window automatic shutter to protect the read/write disk surface, and a write-protect mechanism. In 1000s, prices range between $125 and $200 each, depending on the model. **Panasonic Industrial Co, Computer Components Div**, One Panasonic Way, Secaucus, NJ 07094.

Circle 280

Laser technology delivers disks with 1 Gbyte of long-term storage

At the heart of the LaserDrive 1200 digital optical disk drive are media composed of two 12-in. diameter glass disks. These are coated with a sensitive metallic layer upon which data is written. The LaserDrive 1200 stores up to 1 Gbyte for as long as 10 years on one side of a densely packed disk. The drive has a burst transfer rate of up to 2 Mbytes/s and a 44-Kbyte data buffer that stores interim data. In quantity, the recording disk sells for $265, and the LaserDrive 1200 for $6,600. **Optical Storage International**, PO Box 58063, 3333 Scott Blvd, Santa Clara, CA 95052.

Circle 281

Winchester holds 10 Mbytes and handles PC compatibles

At the heart of the Trustor 10 is a high performance 5¼-in. Winchester disk drive that stores 10 Mbytes of data. Packaged as a self-contained modular unit, it bridges to the host computer via the Xebec S-1410 Winchester controller and is plug compatible with Apple IIs, IBM PCs and PC XTs, as well as various PC compatibles. Price for one unit is $1450. **Datamac**, 432 Lakeside Dr, Sunnyvale, CA 94086.

Circle 282

DATA COMMUNICATIONS

Local area network permits peripheral sharing

Local area networking joins with CAD in the GerberNet 1000. This LAN, developed expressly for the graphics arena, connects multiple Gerber CAD units and plotting systems for peripheral sharing. For example, PC 800 Model 4 CAD units and 3100-based plotting systems can efficiently share common magnetic and punched tape I/O units, printers, and checkplotters. Network interface adapters are used. It supports linking of up to 64 devices over cable lengths reaching 1 km (3280 ft). Coaxial cable is used, and carrier sense multiple access with collision detection is supported. **Gerber Scientific Instrument Co**, PO Box 305, Hartford, CT 06101.

Circle 283

Ethernet server offers low per port connection cost

Linking up to 14 terminals or personal computers, the CS/100-14 acts as a terminal server to link an Ethernet LAN to devices equipped with RS-232-C interfaces. Typical application is as a frontend device for a host computer system. It places no restrictions on the number of terminals or peripheral devices that can be linked to a given Ethernet. Based on a multiple 68000 16/32-bit processor design and an Ethernet chip, the server supports industry standard XNS high level network protocols. Device communication executes via user command interface software. Price is $5400. **Bridge Communications, Inc**, 10440 Bubb Rd, Cupertino, CA 95014.

Circle 284

Protocol converter captures mainframe data and transforms it for PCs

The PA100 Turbo is an IBM PC and XT plug-in board that allows the PC to emulate a variety of 3278 and 3279 terminals. It connects coaxially to a 3274/76 cluster controller supporting BSC or SNA/SDLC environments. Three modes of mainframe access are provided: data capture, file transfer, and terminal emulation. Data capture allows users to easily extract information from existing mainframe applications. Optional host resident software provides a simple menu-driven approach to file transfer. Built-in record and playback feature lets users or the data processing manager create simple one-touch commands for applications. **Avatar Technologies Inc**, 99 South St, Hopkinton, MA 01748.

Circle 285

Protocol analyzer extends performance with software/firmware

Designed to run on the HP 4955A, the HP 18146A decode pack handles CCITT's No. 7 protocol. Option 003 allows translation of the HP 4955A full decode and display capabilities into the Japanese Katakana character set using JIS-8 data code. Further, a remote-testing pack allows an HP 4955A to talk to an HP 4951A or another HP 4955A at a remote location. The HP 4955A with Basic is $17,880; option 003 (JIS-8 data code) is $250; HP 18146A CCITT No. 7 decode pack and HP 18145A remote tester cost $500 each. **Hewlett-Packard Co,** 1820 Embarcadero Rd, Palo Alto, CA 94303. **Circle 286**

Micros talk to mainframes with protocol unit

Apple, Kaypro, Tandy, and IBM PCs link to IBM mainframes with the Hydra II protocol converter. The device enables attachment of ASCII terminals, printers, and microcomputers. Prices are, respectively, $6900 and $9900 for the 8- and 16-port configurations. **Diversified Data Resources, Inc,** 25 Mitchell Blvd, San Rafael, CA 94903. **Circle 287**

Processor can network to ASCII/async RS-232-C computer or peripheral

The IP-3 interface processor has three RS-232-C ports that can attach to virtually any similar device. As many as 64 of the units can be tied together via a twisted/shielded wire bus. Each processor has its own CPU with its operating system stored in firmware. All communication functions can be performed without affecting host CPUs or operating systems. The IP-3 features a 16-Kbyte segmented buffer, password security, and automatic network stats. Price is $1290. **Complexx Systems, Inc,** 4930 Research Dr, Huntsville, AL 35805. **Circle 288**

Workstation doubles as Telex machine

A compact workstation, the PComm, combines full-time communications and computing. It supports Telex, International Telex, teletypewriter exchange (TWX), and other protocols. It also performs microcomputer functions under CP/M, supporting Basic, and running programs such as WordStar, SuperCALC, and dbII. PComm units, available in the fourth quarter of 1984, are priced starting at $3000. **International Digital Electronics Assoc, Inc,** 6 Westchester Plaza, Elmsford, NY 10523. **Circle 289**

Modem employs five microprocessors for fast communication

The WD212-X is a synchronous/asynchronous 1200-bit/s, full-duplex modem. It features Bell 212A, Bell 103, and CCITT X.25 LAPB compatibility for error-free transmission. Self-adaptive program logic is designed into the controlling micros, enabling the modem to adjust automatically to the operating conditions of each telephone line. Direct connect and autodial/answer suit the modem for computer to computer, remote data access, and LAN communication. **Wolfdata, Inc,** 187 Billerica Rd, Chelmsford, MA 01824. **Circle 290**



MPD Solves Computer Designers' Worst Fear.

(*Fear of Memory Loss)

Loss of memory caused by a power failure can be devastating! For example, where vital programs and procedures are stored on CMOS I.C.'s in a computer...

Fortunately, MPD is on guard. Memory Protection Devices' innovative line of state-of-the-art battery holders provide safe, economical and dependable memory back-up systems that can be integrated right into the circuit.

Choice of type is virtually limitless. MPD offers battery holders for low-profile coin cells, stacking coin cells and single or multi-pack cylinder cells...sized to accept Lithium, NiCd, Carbon-Zinc or Alkaline batteries. We even develop custom designs.

And nothing is left to chance. Coin cell and single holders are tough, warp-free Valox® with stainless steel nickel- plated contacts. For multi-cylinder cell holders, material is heavy-duty nylon with spring steel tin-plated contacts. All are designed for PC Board mounting, correct polarity and easy battery installation and removal.

So, if your electronic memory is priceless, remember MPD. Call or write today.



COIN CELLS
- For single Lithium batteries 20-24mm
- Soon stacking holder for 6-volt operation, 23mm and 20-24mm



SINGLE HORIZONTAL
- For Lithium, NiCd, Carbon-Zinc, Alkaline batteries
- PC, fast-on or hardwire connect
- 2/3A, 1/2AA, AA & C.



HORIZONTAL PACKS
- For AA and AAA standard NiCd, Carbon, Carbon-Zinc Alkaline batteries
- 2, 3 & 4 cell series connect.



HORIZONTAL PACKS
- Other styles available, 6" lead wire for connection, male and female 9 volt snap-on terminal, P.C. terminals.

MPD
Creative Solutions to Memory Problems

MEMORY PROTECTION DEVICES INC.
320 Broad Hollow Rd.,
Farmingdale, N.Y. 11735
(516) 454-0340

CIRCLE 117

Low cost modem has intelligent features

Mark X addition to Signalman line is a smart, 300-baud auto-dial/auto-answer modem that works with both tone and pulse dialing. The Mark X runs on most popular software communication packages (eg, ASCII Express). It works manually through a keyboard without computer coding, or to automatically answer and originate calls at 300 bits/s. This unit employs standard RS-232 serial interface with built-in cable. Modem comes with 12-V power supply. Price is $169. **Anchor Automation,** 6913 Valjean Ave, Van Nuys, CA 91406. **Circle 291**

September Preview
Special Report on Computer Systems

Modem exhibits microprogam compatibility

A 1200/300-bit unit, the MultiModem, uses the same command set and functions as the Hayes Smartmodem 1200. It stores up to six 31-digit telephone numbers in its battery-backed memory, detects dial and busy tones, and can continuously redial a busy telephone number until it connects. Modem also features built-in speaker, seven status LEDs and integral test mode. Product is compatible with microcomputer communication packages including CrossTalk, ASCOM, Transend, and ASCII Express. Price is $549. **Multi-Tech Systems, Inc,** 82 Second Ave SE, New Brighton, MN 55112. **Circle 292**

Software analyzes design of communication circuit

Design Kit software allows specialized analysis and optimization of communication circuit design. Available for Tektronix 4051/52/54, Hewlett-Packard series 9816/26/36, and series 500 systems, package includes: RF Design Kit, which optimizes complete system noise figure and intercept point; PLL Design Kit, which optimizes Type-2 second-, third-, and fifth-order loops, and presents general PLL analysis; and Communications Design Kit for dual-loop AGCs, digital filters, and UHF oscillators. **Communications Consulting Corp,** 52 Hillcrest Dr, Upper Saddle River, NJ 07458. **Circle 293**

Small voice-forward unit holds up to 48 hrs of messages

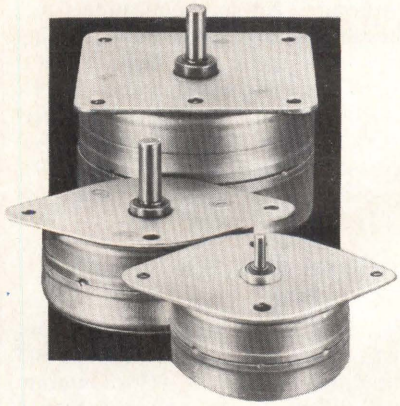
An eight-port voice message exchange system, called the 8S, provides 16 hours of voice storage for PBXs, with optional configuration of up to 48 hrs. It can be integrated with the GTE GTD-4600, InteCom IBX, Northern Telecom SL-100, and other PBX Centrex systems. **VMX, Inc,** 1241 Columbia Dr, Richardson, TX 75081. **Circle 294**

You can count on Synchron® Steppers for performance & economy.

You get prompt delivery too!

In printer, disc drive and sheet feeder applications where the combination of high performance and low pricing are major requirements, be sure to specify Synchron® steppers for full satisfaction.

Up to 46 oz. in. of torque. Up to 5,000 pps. Sizes 19, 23, and 28. Choose from stock models or let us design to your exact specifications.



IMC/HANSEN
a subsidiary of IMC MAGNETICS CORP.
P.O. BOX 23, PRINCETON, INDIANA 47670

Phone 812/385-3415
TLX 278458
HANSEN PRON

CIRCLE 118

Memory cards link MicroPDP/11 systems

Six DEC MicroPDP/11 systems can be connected to a MicroPDP/11 host via the HEX-Q system. This provides up to 28 Mbytes of RAM linked by DMA with transfer rates exceeding 125,000, sixteen-bit words/s. One dual-height DMA card in each computer and two more dual-height cards in the host comprise the HEX-Q system, which is configured in a radial formation. Host and satellite may be separated by as much as 40 ft; a single 40-conductor flat cable running from the host to each satellite is required. Hardware price for a six-satellite configuration is $4500. **Peritek Corp,** 5550 Redwood Rd, Oakland, CA 94619. **Circle 295**

Multiplexer works like X.25 packet switcher

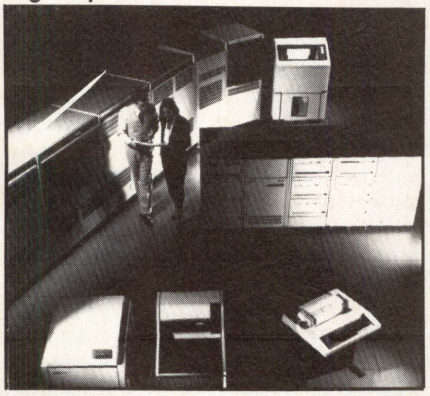
The TP-400 polling statistical multiplexer enables multidrop network configurations to replace existing star networks while using existing host computer software and dumb ASCII terminals. Essentially a closed X.25 packet switcher, the TP-400 employs an 8-bit microprocessor and includes memory and error control software. It is controlled through a port that allows local or remote access by any ASCII display terminal for management control and optimization. **TeleProcessing Products, Inc,** 4565 E Industrial St, Bldg 4, Simi Valley, CA 93063. **Circle 296**

Telephone unit monitors remote sites from central location

The remote telephone monitor (RTM) consists of two units. The master unit is at a central site while the slave unit is at a remote site. The slave unit, connected to the remote phone line, samples analog signals at the Nyquist rate, storing the data in 16 Kbytes of RAM. Using error correcting algorithms, the data is transmitted to the master and stored in the master's memory. The remote signal is then restored, allowing it to be analyzed by measurement equipment or connected to a modem. Price is $1500 per unit. **RAD Computers, Ltd,** 871 Seventh Ave, New York, NY 10019. **Circle 297**

COMPUTERS

Mainframe upgrade means higher performance



Higher capacity versions of DECsystem-10 and DECsystem-20 models exhibit a 20-percent performance increase over current models. These mainframes, the DECsystem-1095 and DECsystem-2065, use a 36-bit CPU, a cache/pager with 18 Kbytes of memory, 176-Mbyte disk drive, and 16 asynchronous-line front ends. Fully software-compatible with existing family models, system prices start at $395,000 for the DECsystem-1095, and $355,000 for DECsystem-2065. Upgrade kits are priced at $40,000. **Digital Equipment Corp,** 146 Main St, Maynard, MA 01754. **Circle 298**

Talk to the editor

Have you written to the editor lately? We're waiting to hear from you.

Demand-paged virtual memory resides in Unix-based system

Based on the Unix System V operating system, the PowerStation 2000 runs the universal time-sharing executive (UTX) operating environment. The unit supports up to eight users and uses the 32-bit 68010 operating at 10 MHz with no wait states. Memory is provided in 0.5-Mbyte increments with expansion to 2 Mbytes. It supports the PT100 intelligent terminal for 307-kbit/s communication. A user-friendly operating system, the UTX/2000 features window manager and key prompt shell program. Price is $8995. **Gould Inc, Computer Systems Div,** 6901 W Sunrise Blvd, PO Box 9148, Fort Lauderdale, FL 33310. **Circle 299**

Supermicro offers 8-in. Winchesters and increased capacity

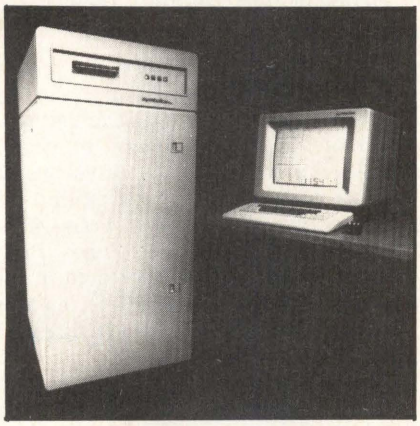
The Universe 68/137 incorporates compact 8-in. Winchester disk drives with 120 Mbytes of formatted capacity. Sporting a 45-Mbyte backup tape cartridge, this system is available with 512 Kbytes of parity checking memory. The system is based on the 68000 and uses the 32-bit VERSAbus for I/O. Like previous Universe 68 family members, it maintains 32-bit main memories and 32-bit cache memory. Universe 68/137 runs under Unos, a proprietary, Unix-compatible realtime operating system, or Un/System V, a Unix System V implementation. The 512-Kbyte parity checking memory version is $26,900, while a model with 1 Mbyte of error checking and correcting memory is set at $29,650. **Charles River Data Systems,** 983 Concord St, Framingham, MA 01701. **Circle 300**

Multitasking systems handle standalone or network chores



Multi-user, multitasking Zeus and Jupiter work centers are built around the LSI 11/23 and 73. Using DEC-compatible operating systems, languages, and application software, the desktop machines can operate in networks or as standalones. A card cage with 10 slots is offered, as well as 512-Kbyte RAM and 22-bit addressing for direct access of up to 4 Mbytes of memory. The devices accommodate a 20.8-Mbyte Winchester subsystem, a 1-Mbyte floppy disk subsystem, four RS-232 serial ports, a CRT terminal, and a detached solid state keyboard. Price is less than $10,000. **Spectra Systems, Inc,** 2754 Compass Dr, Grand Junction, CO 81501. **Circle 301**

Symbolic processor offers proprietary tagged memory architecture



The 3670 computer features a backplane with 14 optional expansion slots, an increase of seven over the 3600 processor. It can support as much as 30 Mbytes of memory with a 1-Gbyte virtual memory complement. The FPA option enhances performance by executing floating point operations in parallel with data-type checking. The output processor speeds video operations 2.5 times by implementing a buffered, pipelined pixel memory system. A cache between the processor and video memory and a FIFO between video memory and the display act to enhance graphics performance. Base price is $84,500. **Symbolics, Inc**, 4 Cambridge Center, Cambridge, MA 02142.
Circle 302

Raster display and graphics systems feature concurrent processing

The LEX 90 series devices use a display controller based on a 56-bit word microprocessor, and writes vectors at 600 ns per pixel in any direction. Multiple read/write and arithmetic functions are handled concurrently. Available in four upgradable versions, the LEX 90/35 Model 2 is a dual-resolution graphics system that allows simultaneous display of both a 640- x 512- x 8-pixel image and a 1280 x 1024 x 4 overlay on one screen. The LEX 90/35 Model 3 is a high resolution graphics system supporting eight planes of 1280 x 1024 display memory at 60-Hz noninterlaced refresh. Model 2 prices range from $13,925 for four-plane to $22,925 for true color system. Model 3 is $19,850 for an eight-plane configuration. **Lexidata**, 755 Middlesex Tpke, Billerica, MA 01865.
Circle 303

Array processor family enhances reliability

The AP400 series for DEC, Data General, and Hewlett-Packard computers has a reduced parts count for higher reliability. The prior family exhibited an MTBF in excess of 8500 hrs—the enhanced systems expect to reach an MTBF in excess of 12,000 hrs. Affected products are the DEC compatible 421/426 (chassis version) and 471/476 (card set), the Data General compatible 423 (chassis version) and 473 (card set), and the HP compatible 424 (chassis version). All family members are hardware and software compatible with their predecessors and are available 30 days ARO. **Analogic Corp**, 8 Centennial Dr, Centennial Industrial Park, Peabody, MA 01961.
Circle 304

Single-board computer sports mainframe capabilities

A complete 32-bit computer on a single board, the Python/32, includes CPU with true 32-bit instruction set, arithmetic and data paths, 4 Mbytes of dual-ported error correcting RAM, a storage module disk, streaming tape controllers, and a 64-port serial multiplexer. It measures 16.75 x 25 in. (42.5 x 63.5 cm), and supports Unix. Based on the NS32032, the Python/32 features hardware floating point and onboard DMA that handles high speed data transfer between system components. An Ethernet link and 16 Kbytes of user-programmable ROM are also featured. Quantity price is under $30,000. **General Robotics Corp**, 57 N Main St, Hartford, WI 53027. **Circle 305**

Station based on the 80286 links to the ARCnet

Based on the 80286 microprocessor, the Vista-Station-84 speeds batch processing and features multitasking capabilities. The processor performs functions such as word processing, electronic mail, and windowing within an ARC local area network. A fiber optic hub enables users to extend the distance between standard ARC network hubs supporting the Vista-Station-84 and other ARCnet-compatible units. Maximum distance between hubs can be increased from 2000 to 4000 ft. A low profile ergonomic keyboard is also featured. **Datapoint Corp**, 9725 Datapoint Dr, San Antonio, TX 78284.
Circle 306

INTERFACE

Controllers drive color graphics using high speed bit slice

Omega 300 and 305 color graphics controllers are four bit planes deep. The four bit planes yield 16 display colors from a palette of 16.7 million and feature a high speed bit-slice bipolar processor that is faster than MOS microprocessor-based systems. In addition, pan, zoom, and FLASH-fill give the user significant performance enhancements. For example, the Omega 300 FLASH-fill rate hits 16 million pixels/s, and the vector drawing rate is 1 million pixels/s. The Omega 300 is $5950—the Omega 305, $6950. **Metheus Corp**, PO Box 1049, Hillsboro, OR 97123.
Circle 307

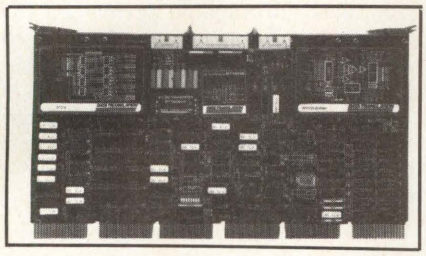
Controller provides graphics functions in CMOS VLSI

The ACRTC features high resolution, logical X-Y address translation, windowing, clipping, and fast graphic drawing. The board connects to a variety of monochrome and color CRTs and introduces fully programmable CRT timing signals, interlaced and two noninterlaced scanning modes, and synchronization with multiple ACRTCs or other video signals. The HD63484 combines three onchip processors (drawing, display control, and timing) that operate in parallel using pipeline control techniques. It will support a 2-Mbyte bit-mapped graphics address. **Hitachi America, Ltd**, 1800 Bering Dr, San Jose, CA 95112.
Circle 308

Interface upgrades older printers to daisy wheel performance levels

Printerface models DT150 and 151A attach to printers to provide standard interface configurations including RS-232 serial, Centronics parallel, and IEEE 488. They provide automatic bidirectional printing, proportional spacing, boldfacing, autocentering, variable pitch, and self-test mode. These intelligent interfaces are designed for the Diablo Hytype I and II, DEC LQP-01, and Xerox printers. **Kuzara International**, 770 Vickers St, San Diego, CA 92111.
Circle 309

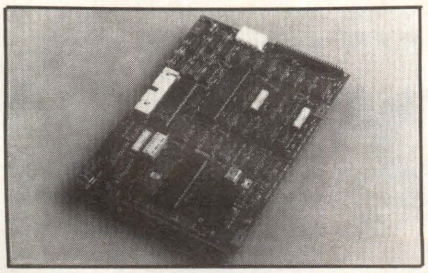
Unibus hardware/software combo helps VAX data acquisition



The DT1771 series consists of analog I/O boards and realtime support software for the Unibus. Using a 16-bit, bit-slice processor operating at 8 MIPS, the series features a 64-Kbyte memory buffer and covers eight modes of analog I/O: continuous data acquisition, countdown data acquisition, pre- and post-event sampling, histogramming, signal averaging, continuous D-A, and D-A with external initiation. Support is available for the RSX-11M, RSX-11M Plus (RSXLIB), VAX/VMS (VMSLIB), and VAX Unix 4.2 BSD (UNIXLIB) operating systems. The DT1771 is priced at $7870. Initial license fee is $1995 for RSXLIB and $3295 for VMSLIB and UNIXLIB. **Data Translation Inc,** 100 Locke Dr, Marlboro, MA 01752.
Circle 310

Tape controllers make SCSI bus connections

Titleist tape controllers connect the Cipher 540 (a quarter-inch streaming tape drive) with the SCSI. The Titleist can interface through the SCSI bus to CPUs, including the Q-bus and Unibus, the Multibus, and IBM PC XT bus structures. When interfacing a Cipher 540, the controller uses the QIC-24 media format, which provides nine-track sequential, serpentine recording with 8000-bit/in. density at 90-bit/in. speed. The Titleist supports full SCSI arbitration for up to eight host and/or peripheral devices. It also offers disconnect/reconnect. Product lists at $435. **Emulex Corp,** 3545 Harbor Blvd, PO Box 6725, Cost Mesa, CA 92626.



Circle 311

Instrument control system is based on IBM PC

Plus500 instrument control systems interface the IEEE 488 bus to the IBM PC. Up to 15 instruments can be simultaneously managed from a single keyboard. This system allows automatic serial polling and adjustable time-outs. Plus500 software, which operates as an extension of Basic, is included. **Keithly DAS,** 349 Congress St, Boston, MA 02210.
Circle 312

4913

COLORS YOU CAN COUNT ON

1 COUNT ON COLOR. Without even realizing it, we are constantly responding to color stimuli everyday, everywhere. Color attracts attention. Color emphasizes. It distinguishes, accents and highlights. In essence, color communicates and that really counts. And color is counting even more as needs increase for information to be comprehended faster, more accurately and with greater retention.

2 COUNT ON TC1040. For effective color communications the TC1040 has a unique 4,913 color palette — a powerful vocabulary you can count on. Our outdoor ink-jet system prints up to 120 dots per inch, delivering crisp and clear type fonts, graphics and images on a variety of media. Push botton console, flexible media handler and snap-in disposable cartridges make operation as easy as 1-2-3. No complicated procedures, no elaborate set-ups, no messy pens or ribbons to replace. And the TC1040 is fast — full image-seeking and bidirectional capabilities will print an 8½" × 11" page in just one and a half minutes.

3 COUNT ON PRINTACOLOR. Since our introduction of the first desk-top color ink-jet printer in 1980, we have advanced ink jet color technology to a high level of refinement, sophistication and readability. You can count on our experience for putting quality color into your business graphics, audio-visual presentations, process control operations, scientific data analysis, medical imaging. Whatever your application, we're ready to put our unique 4,913 color communicators into it. For details, write or call: Printacolor Corporation, 6040 Northbelt Drive, Norcross, Georgia 30071; 404 · 448-2675.



PrintaColor

CIRCLE 119

Digital to analog converters have built-in micro interface

The DAC708 and 709 are 16-bit devices that combine a DAC700 D-A converter with a dense logic CMOS gate array in a 24-pin double wide DIP. Maximum linearity error is ±0.003 percent FSR. The devices are monotonic to 14 bits over their specified temperature ranges. Gain drift is ±10 ppm/°C typical. Both voltage out and current models are available. The micro interface is composed of two 8-bit input registers and one 16-bit D-A register, each with its own enable line. The D-A converters are priced from $44 in 25 plus quantities. **Burr-Brown Corp, Data Products Div,** PO Box 11400, Tucson, AZ 85734. **Circle 313**

Multiprocessor system integration gained by host-independent interface

Asynchronous bit parallel communication between two independent processors is possible with the TMS9650 multiprocessor interface. The TMS9650 is host-independent. It connects 8-, 16-, or 32-bit wide microprocessors capable of interfacing to standard memory or peripheral devices. Consisting of eight programmable 8-bit registers at each of its two ports, the device provides access to 256 bytes of onchip RAM used to buffer data transferred between ports. It requires little software overhead and supplies internal arbitration logic to resolve RAM access conflicts. These 5-V, 40-pin DIP units cost $13.20 each in 10,000-piece quantities. **Texas Instruments,** PO Box 809066, Dallas, TX 75240.
Circle 314

Analyzer handles IEEE 488 bus interface chores

Designed as an accessory for Erbtec's programmable interface (EPI) board, the Bus Analyzer Module allows bus implementation without development of unique system analysis tools. Two major operational modes are available: bus analyzer mode, in which data and control signal information present on the IEEE 488 bus is displayed; and addressed device mode, in which the module displays signal data present on the EPI board's two 8-bit I/O ports. **Erbtec Engineering, Inc,** 5680 Valmont Rd, Boulder, CO 80301.
Circle 315

Overlay controller can mix graphics and video

A graphics overlay/controller board allows PC-created graphics to unite with images from video sources. Video disk or tape player output can also be controlled. The board mixes the output from two graphics cards, and output is available in two formats. The board will lock onto an incoming video source, or it can supply a composite sync signal to a player. Available now for the IBM PC and PC lookalikes, the price is $450. **IEV Corp,** 254 W 4th S, Salt Lake City, UT 84101.
Circle 316

RELIABILITY–REMOVABILITY

PROVEN PERFORMANCE WORLDWIDE IN DISK DRIVES AND DRUM MEMORIES

Behind every one of VRC's rugged disk and drum memory products is a record of reliability known and unchallenged for a quarter of a century. VRC drives are designed to operate with a high MTBF and tolerance to thermal extremes (0° to 55°C), shock and vibration. VRC has two new cartridge drives that make reliability compatible with removability — the 8520 and the 8010. Both offer a removable 11-megabyte, 8-inch ANSI cartridge plus the security of off-line data storage and back-up. The 8520 also has an 11-megabyte fixed disk in the same compact package. The VRC 8000 series drives incorporate non-contact, high-flying heads that never touch the media in stop, start or transit modes. A proprietary embedded servo head positioning system guarantees cartridge interchangeability and eliminates head



alignment problems. The result is a removable media product with high tolerance to environmental extremes.

VRC's head-per-track drives represent the state-of-the-art in drum memory evolution. For applications such as telecommunications, process control and computer aided manufacturing, VRC offers a 25,000 hr. MTBF and a ten-year design lifetime for both the 9.6-megabyte 4040 and the 4.7-megabyte 4016. They can replace every early or current fixed-head device that has a digital interface. All systems are supported with interfaces, controllers and power supplies and extensive documentation.

For more information call or write:

VERMONT RESEARCH CORPORATION
Precision Park
North Springfield, VT 05150
Tel: 802/886-2256
TWX: 710/363-6533

In Europe:

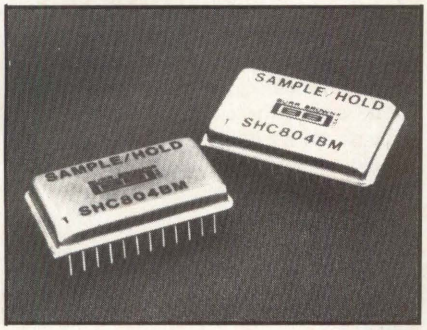
VERMONT RESEARCH LIMITED
Cleeve Road
Leatherhead, Surrey
England KT227NB
Tel: 0372-376221
TLX: 895 4667

VRC

Memory Products for Systems That Can't Stand Failure

CIRCLE 120

Data acquisition sample/hold amps display fast settling times

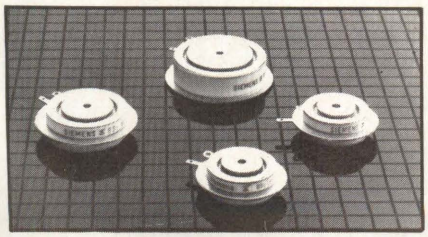


The SHC803 and SHC804 sample/hold amplifiers exhibit a maximum acquisition time of 350 ns for settling to within ±0.012 percent of 10 V (±1.2 mV). Apt for fast 12-bit data acquisition systems, these units operate as unity-gain inverters in the sample mode. Their signal bandwidth equals 15 MHz, with an input voltage range of ±10 V. Output drive current is ±50 mA. The SHC803's uncommitted frontend buffer amp distinguishes it from the 804. Power supply requirements are ±15V and 5 V. In 100s, the SHC803 is priced from $117; the SHC804 from $105. **Burr-Brown**, Box 11400, Tucson, AZ 85734.

Circle 317

Hockey puck inverter SCRs offer wide voltage ratings

Minimum trigger currents for BST-61 inverter SCRs reach 100 mA at a junction temperature of 140°C, with voltage ratings running from 200 to 600 V. These "hockey puck" inverters come in ½-in. high, compression-bonded ceramic disk packages that range in diameter from 37 to 50 mm. Useful in uninterruptible power supply applications, reverse current fall times are as low as 15 μs at maximum on-state current levels up to 1730 A. **Siemens Components, Inc, Colorado Components Div,** 800 Hoyt St, Broomfield, CO 80020.



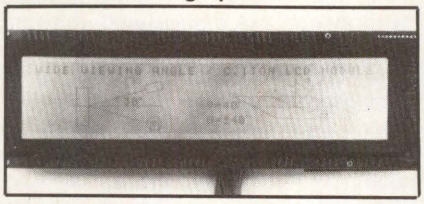
Circle 318

Connectors easily snap to shorter lengths

Standard 36-single and 72-position dual row strips of insulation displacement connectors save wasted board space and pin costs by getting the exact size needed. Cable assemblies are designed to mate with 0.025-in. square posts to form connector sets, board/board and on/board interconnects, jumpers, and test points. Contacts are beryllium copper, gold plated. Cable is 28 awg with PVC insulation on 0.100-in. centers for single row, 0.050-in. dual row. Priced from $0.28 in 100-piece quantity. **Sametec, Inc,** PO Box 1147, New Albany, IN 47150.

Circle 319

Large-area LCD modules provide dot-addressable graphics



High contrast displays gain dot-addressable graphics capabilities with two LCD modules—the CG4801280D and the CG6401282D. With a 1.5-in. (38.1-mm) thickness, voltage requirements of 5 V and −12 Vdc, the modules are apt for battery operated, portable applications. Units employ advanced twisted nematic-type LCD technology, and wide viewing angle CMOS drivers and controllers. Interface is CMOS and TTL compatible. The CG4801280D and CG6401282D, in quantities of 1000, cost $200 and $220, respectively. **C. Itoh Electronics, Inc**, 5301 Beethoven St, Los Angeles, CA 90066.

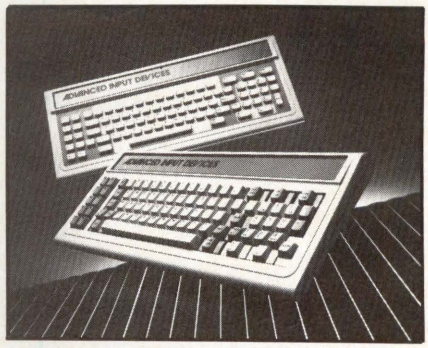
Circle 320

Flat electroluminescent display fits into portables

The MDM 512.256-11 module has 512 x 256 picture elements capable of displaying 25 lines of 80-char text or high resolution graphics. High picture quality stems from the subwavelength thin light emitting EL phosphor layer. The solid state panel and board with required high voltage drivers are assembled into a 0.37-in. thin rugged package, which connects to controlling circuitry over a flat cable. In 1000s, the price is $700. **Finlux, Inc,** 20395 Pacifica Dr, Cupertino, CA 95014.

Circle 321

Full-travel keyboard features conventional DIN truncated keytops

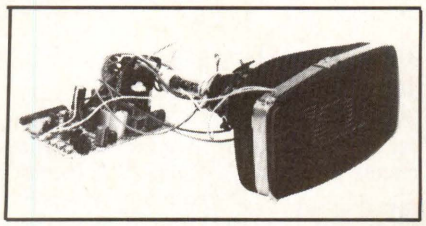


The 83-key Ergokey EKT complements the EKI version, which offers international rectangular style keys for quick-change graphics overlays. Features are DIN profile, lightweight construction, choice of 2- or 3-oz operating forces, IBM tactile feel, and an elastomer one-piece switch mat. The high life mat (greater than 60 million cycles/switch) provides spillproof and ESD protection to circuitry. The standard electrical output is a DIN connector with a 6-ft shielded coil cord that is IBM plug compatible. Cost is less than $50 each. **Advanced Input Devices,** W 250 AID Dr, Coeur d'Arlene, ID 83814.

Circle 322

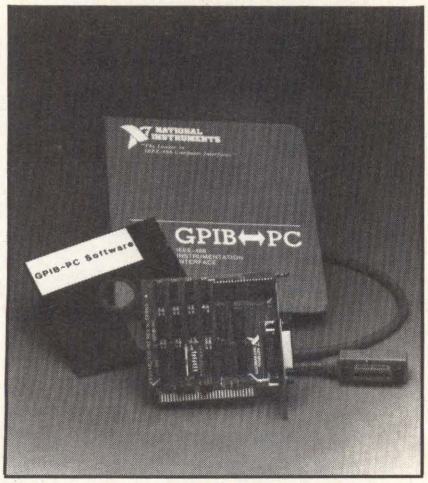
Data display presents information on low profile screen

The Model 9DD960 has a 9-in. horizontal screen configuration, but the vertical height is condensed to 5 in. This maintains a fully usable display area of 3 x 8 in. Available with the standard P4 phosphor, it comes in kit form for installation flexibility. The unit features 700-line resolution, a video bandwidth of 25 MHz, and horizontal scanning frequency of 15.75 to 22 kHz. Options include choice of phosphors and glare treatments. **Audiotronics Corp,** 7428 Bellaire Ave, North Hollywood, CA 91605.



Circle 323

Software converts computer to instrumentation controller



Designed for GPIB PC interface, the instrumentation software provides a set of primitive GPIB commands for manipulating data management and handshake lines. High level commands are also available and can be mixed with primitives in the same program. Functions may be used interactively from a terminal using the control program. Applications can be written in interpretive or compiled Basic, Fortran, Pascal, C, or assembler. The software is priced at $75. **National Instruments,** 12109 Technology Blvd, Austin, TX 78727.
Circle 324

Package gains powerful graphics and plotting for IBM PC

The CALCOMP-compatible SKYGRAF software package is written in Fortran. It converts the PC into a powerful graphics and plotting tool. SKYGRAF subroutines are device independent and interface to many standard printers. Libraries implement parts of two industry graphic standards—the GKS and Core—and a simple interface allows users to easily display complex data and charts. **Sky Computers, Inc,** Foot of John St, Lowell, MA 01852.
Circle 325

Operating system delivers latest Unix enhancements

A Unix derivative operating system, XELOS, incorporates the latest enhancements of AT&T's 5.2 release. It includes the standard AT&T command interpreter shell, plus the C shell from the University of California at Berkeley. The C and Fortran-77 compilers, a 3200 assembler, and a symbolic debugger are included. Unix multi-user, multitasking facilities are present, as is a hierarchical file system with support for flexible access protection and shared files, and a record locking enhancement. Prices range from $1500 for an eight-user license on the Model 3205 to $30,000 for 65-plus users on the 3250XP. **Perkin-Elmer Data Systems Group,** 2 Crescent Pl, Oceanport, NJ 07757.
Circle 326

Graphics software runs on popular microcomputers

Micro Template represents a PC version of Template interactive graphics software. It provides two-dimensional drawing, geometric, text, input, and polygon primitives. A special fill algorithm increases power and speed for polygon fills. Versions for the IBM PC and XT as well as for the TI Professional Computer, are available, with versions for PC-compatibles due soon. Micro Template price is $300. **Megatek Corp,** 9605 Scranton Rd, San Diego, CA 92121.
Circle 327

Basic to C translation guidance performed for Unix applications

The Main Selector Menu automatically loads S-Tran and guides users through each phase of the Basic to C translator process for Unix and Unix-like applications. Learning and executing all commands required by Unix for Basic translation is no longer necessary. The Main Selector Menu also assists transfer of Basic programs and files to a Unix-based system. This package can also list errors occurring during translation or compilation. **Software Manufacturers Inc,** 20720 S Leapwood Ave, Carson, CA 90746.
Circle 328

Software allows user signal processing without programming

Signal processing software has been adapted to run on DEC's Micro-PDP/11 under the RSX or RT-11 operating system. Designated "Workstation ILS," the package allows users to process and analyze signals without programming. The system provides three-dimensional displays, digital filtering, signal editing, pattern analysis and spectral density as well as modeling, correlation, convolution, and coherence. Workstation ILS consists of 80 integrated programs composed of 50,000 lines of Fortran code. Price ranges from $4500 to $12,500 depending on the operating system and distribution form—source or binary. **Signal Technology, Inc,** 5951 Encina Rd, Goleta, CA 93117.
Circle 329

System 34/36 file access extends to the PC

Installation of DB/LINK modules allows users to access, select, and process information from an IBM System 34/36 and then transfer this data to an IBM PC or PC XT. Further, DB/LINK allows access to multiple file data bases so data can be selected and sorted, and eliminates the need for transferring whole files. This package is an add-on module to ExecuTrieve, an information management system for the System 34/36. DB/LINK gives users the choice of formatting transferred data for use with ExecuTrieve software on the PC or for other micro software programs. **On-Line Software International, Inc,** Fort Lee Executive Pk, Two Executive Dr, Fort Lee, NJ 07024.
Circle 330

Compiler handles ISO Pascal for 80286 applications

Pascal-286, a high level language translator, aids iAPX 286 application program development. This compiler conforms to ISO Pascal standards and is upwardly compatible with programs written in Pascal-86. Pascal-286 produces relatively compact object code (characterized by short execution time) because it makes efficient use of iAPX 286 instructions for string handling, coprocessor numerics and subroutine linkage. A common set of well-defined, well-documented operating system interfaces allows easy transition from the development host to the target system. Labeled the iMDX-324, this Pascal-286 software package is available now for $3900. **Intel Corp,** 3065 Bowers Ave, Santa Clara, CA 95051.
Circle 331

IF YOU THINK
THAT SIGNETICS
JUST CRANKS OUT A BUNCH
OF JELLYBEAN PRODUCTS,
IT'S TIME YOU TURNED
THE PAGE.

OUR VMEbus HAS A SIGNIFICANT EDGE.

First, it has Signetics and Philips behind it. That puts our leading-edge VMEbus in a world class by itself.

Philips is a $17 billion multinational electronics corporation with vast R & D resources and broad applications experience. Engineers and programmers from both companies are working on VMEbus systems around the world. This gives us an edge in refining and expanding our VMEbus board family.

Second, we give you a really competitive edge with our Eurocard connection, one of many quality features we have on board.

It's much more reliable than the standard edge card, because it self-seals to keep out dirt. As you know, a dirty edge connector can shut down an entire assembly line.

Old world money. New world family.

Our VMEbus boards are built around Signetics-developed VLSI products in the 68000 family. These include 15 VLSI communications and control peripheral chips in addition to three CPUs.

So you have a clear advantage in architecture and performance. As well as a migration path from 8 to 16 to 32 bits.

The VMEbus line starts here.

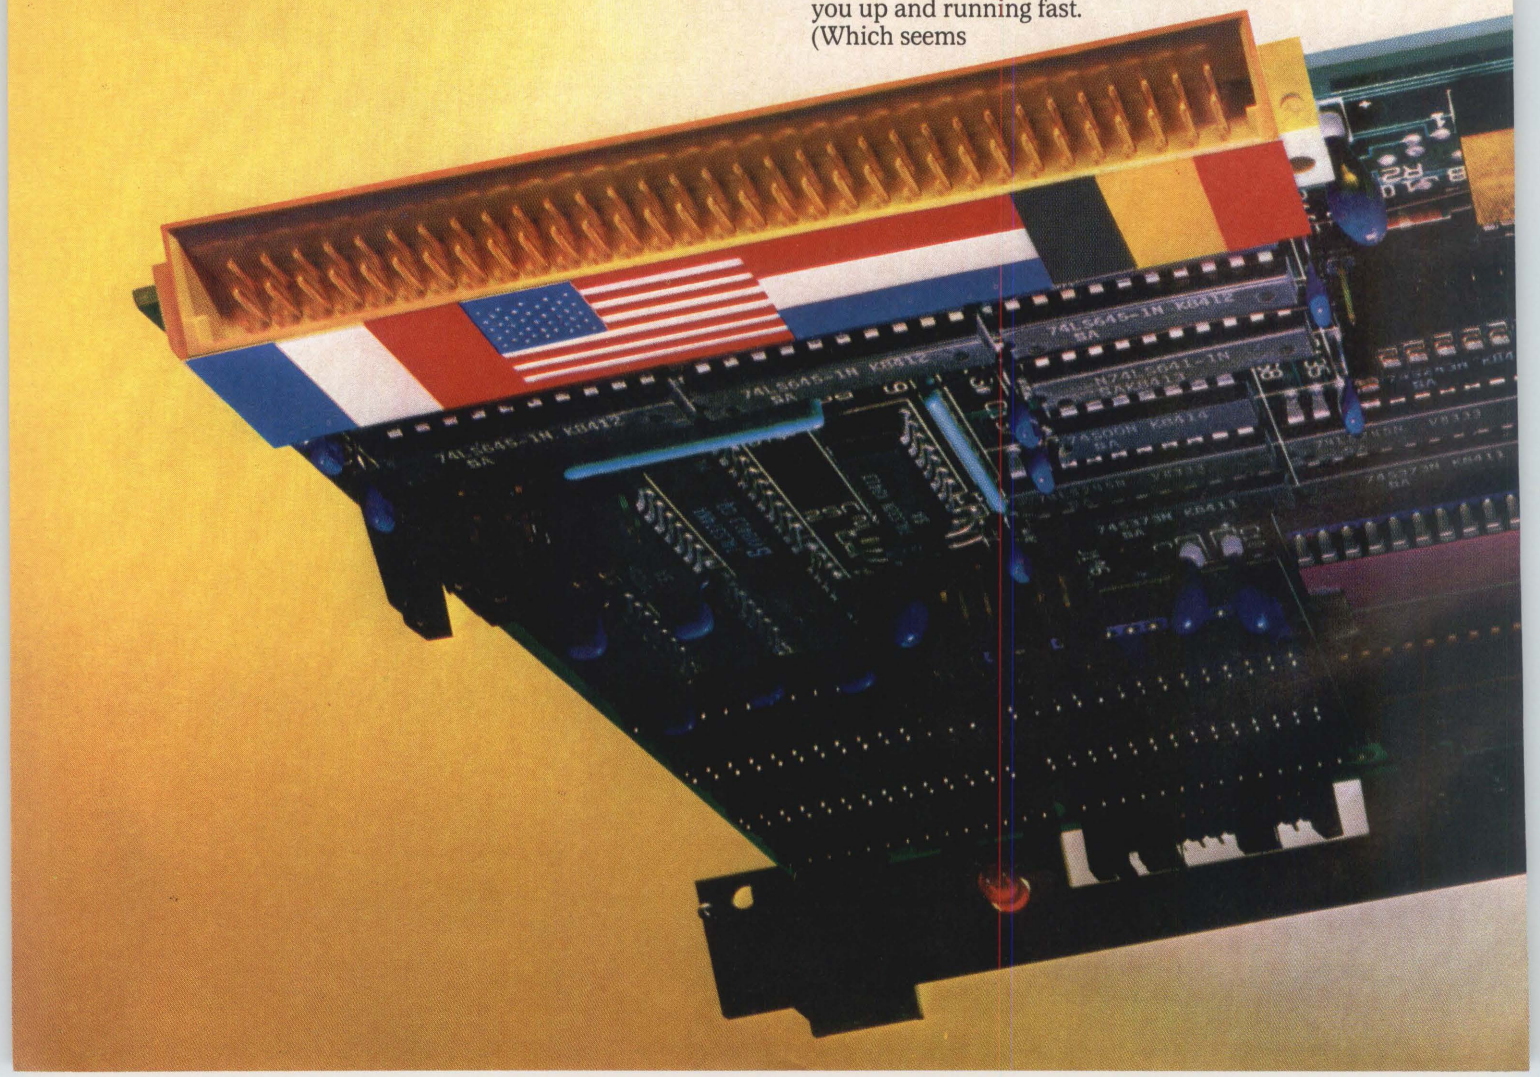
We have nine different boards to choose from, with more to come. The present line-up includes:

SVME 2000 Series	8 and 10 MHz CPU boards
SVME 3000 Series	RAM/ROM/EPROM boards
SVME 4000 Series	Disk Controller boards
SVME 5100 Series	Data Communications boards

Later this year, you'll see new products in our highly compatible family. All of them will reflect our commitment to increasingly sophisticated VLSI.

With every introduction, we're packing more and more performance and functions onto fewer boards. And they all run the same software.

Support is another plus. It's designed to get you up and running fast. (Which seems



appropriate for a bus system that runs two to five times faster than the competition.)

We offer the pSOS-68K Real-Time Executive, together with a pROBE-68K debugger. Cross compilers and macro assembler. And our User Work Station, which interacts with existing computers for hardware/software integration and software debugging. If you need more help, our Field Application Engineers are on call.

An Evaluation Kit that tells all.

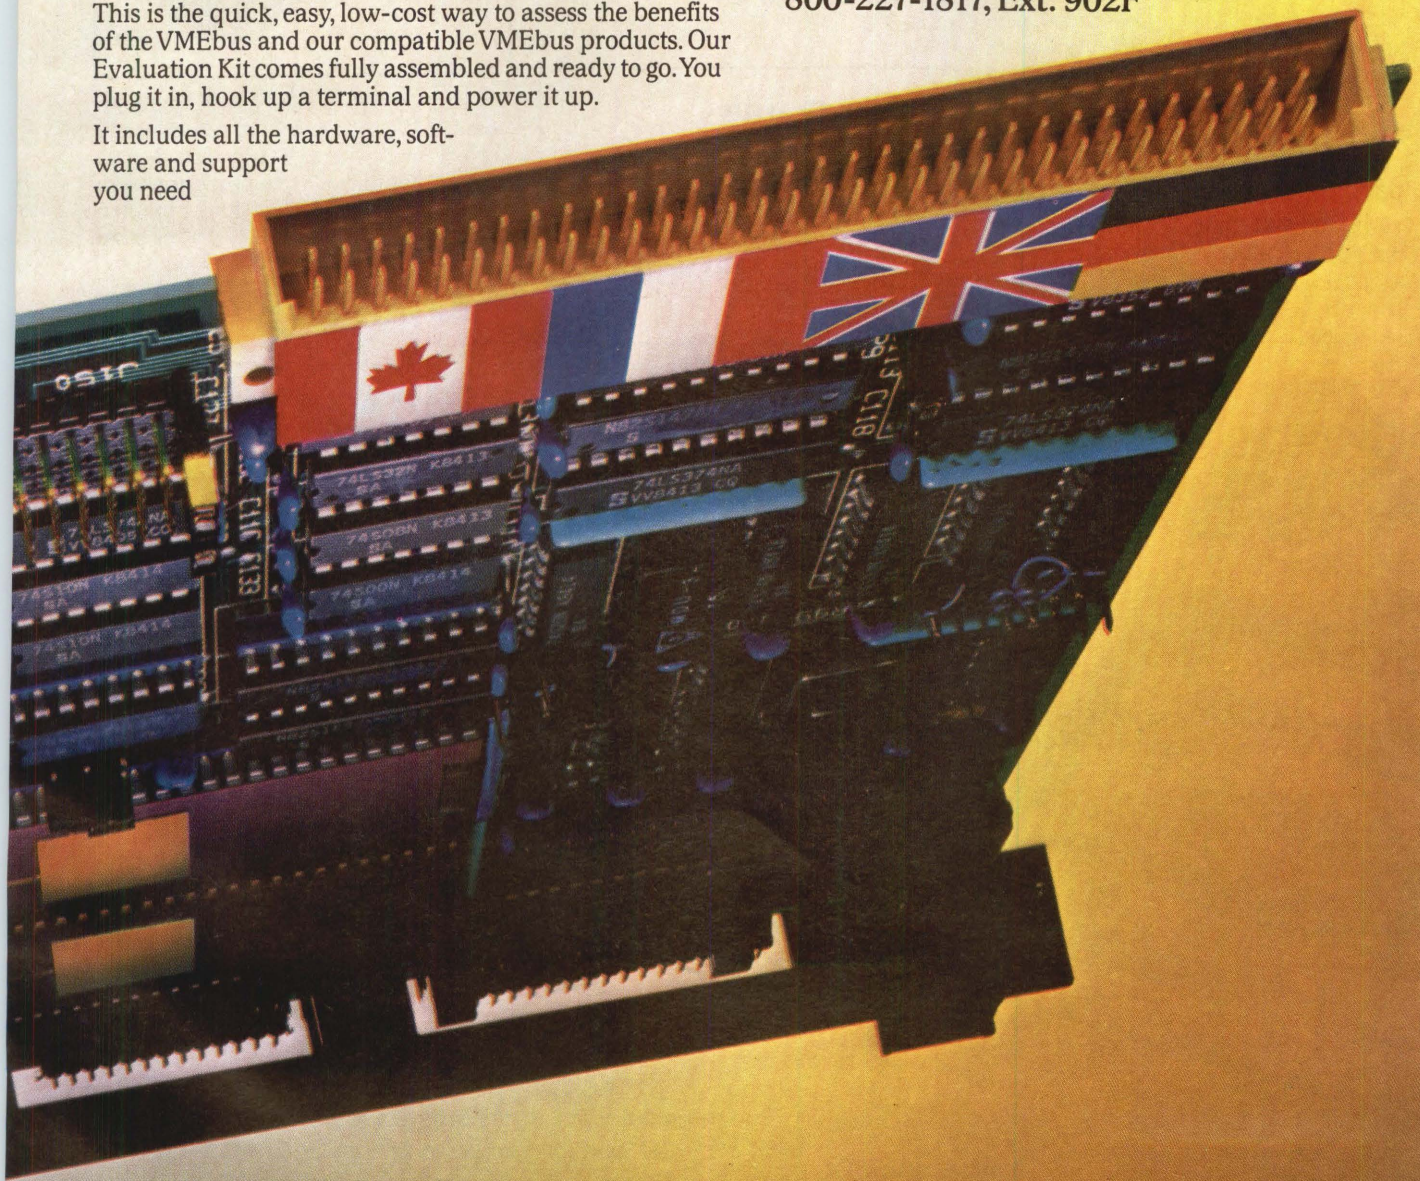
This is the quick, easy, low-cost way to assess the benefits of the VMEbus and our compatible VMEbus products. Our Evaluation Kit comes fully assembled and ready to go. You plug it in, hook up a terminal and power it up.

It includes all the hardware, software and support you need for a thorough evaluation. Afterwards, you can expand it into an economical resident development system.

With a full line of VMEbus boards at your local Signetics distributors, you can get a jump on the rest of the world right now. Contact them for details, including information on the Evaluation Kit. Or call us toll free for literature and the phone number of your nearest Signetics representative.

You'll soon see why our VMEbus gives you such a clear edge.

800-227-1817, Ext. 902F



VLSI from Signetics

a subsidiary of U.S. Philips Corporation

Package supports highly intelligent word processing

Word processing for Unix 68000-based multi-user systems is made possible by CrystalWriter, which offers full editing capabilities found in advanced word processors while meeting complex formatting requirements and providing ease of use. CrystalWriter is object oriented—users can specify the type of document desired and the system automatically adjusts to that document-type's parameters. The package coordinates multiple terminals and printers. Unit price is $1000, with quantity discounts available. **Syntactics Corp**, 3333 Bowers Ave, Suite 145, Santa Clara, CA 95051.
Circle 332

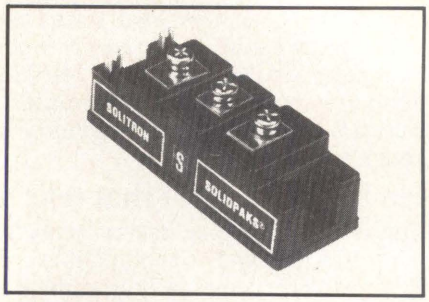
Database management for NCR Tower follows Unix format

A Unix-based DBM system runs on the NCR Tower 1632. Known as Unify, it is a fourth-generation Unix DBMS with multiple access methods and multiprocess architecture. Unify offers interactive data entry, SQL 2 (a relational query language), a report writer, a database update feature, and database load/unload. Full-screen, two-dimensional development tools and a menu-oriented user interface enhance operation. The Unify developer system is priced at $1995, including documentation. The Unify runtime system is priced at $995. **NCR Corp**, 1700 S Patterson Blvd, Dayton, OH 45479.
Circle 333

Like to write?
The editors invite you to write technical articles for Computer Design. For a free copy of our Author's Guide, circle **503** *on the Reader Inquiry Card.*

PACKAGING & POWER

Dual-isolated power modules handle three amp ratings



The SOLIDPAK Series power modules come in 30-, 50-, and 70-A versions, and feature two isolated Darlington devices with an output diode that provides reverse voltage protection. Versions, marked SPK/I-A30, A50 and A70, are priced from $28 each in 100 or more quantities. **Solitron Devices, Inc**, 1177 Blue Heron Blvd, Riviera Beach, FL 33404.
Circle 334

PACKAGED POWER®



POWERING YOU THROUGH THE 80'S... AND BEYOND!

NEW
from the pros in Packaged Power.
A 116-page catalog with complete specs on more than 700 standard power supplies.

AC/DC linears and switchers. DC/DC converters.
The broadest line in the business. It's all here, including prices.
And it's FREE! Send for your personal copy or **Call (305) 974-2442.**

CIRCLE 121

POWER PRODUCTS DIV.
POMPANO BEACH, FLORIDA

COMPOWER DIV.
SAN JOSE, CALIFORNIA

STEVENS-ARNOLD DIV.
SOUTH BOSTON, MASSACHUSETTS

POWER PRODUCTS LTD.
REPUBLIC OF IRELAND

COMPUTER PRODUCTS™
POWER CONVERSION
2981 S.W. 14 St. • Pompano Beach, FL 33069

WHERE THERE'S A WILL, THERE'S VTI.



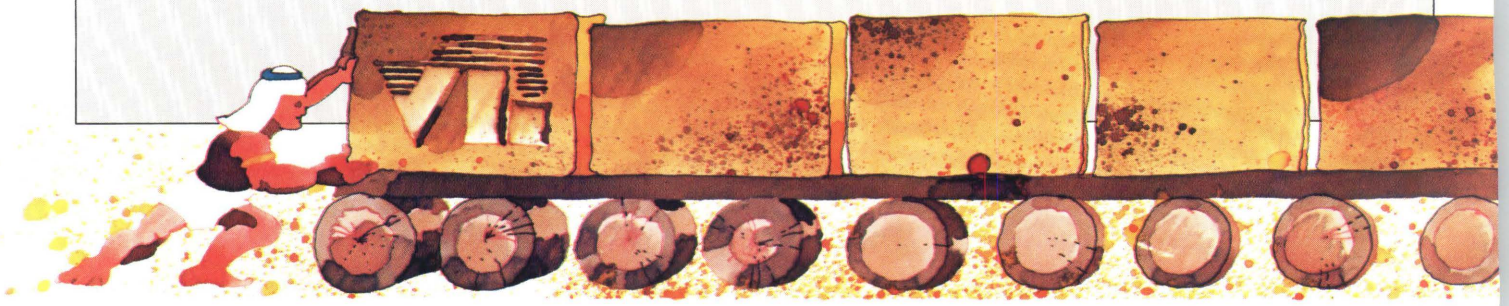
VTI HELPED MINDSET GRAPHICS AT

If the ancient Egyptians hadn't pushed technology to its limit 5000 years ago, one of the seven wonders of the world might never have been realized. The same determination holds true today at VTI. That's the reason Mindset came to us to help design their innovative new personal computer.



The Mindset Personal Computer. VTI's chips enable it to offer $50,000 graphics capability for $1,800.

Mindset envisioned an IBM compatible machine with graphic performance capabilities only available on large, expensive graphic systems. But they needed some expert assistance to make sure the world wouldn't end up with yet another IBM PC clone. The only company able to provide Mindset with all the resources they needed for the immense task at hand was VTI.



DELIVER MAINFRAME A MICRO PRICE.

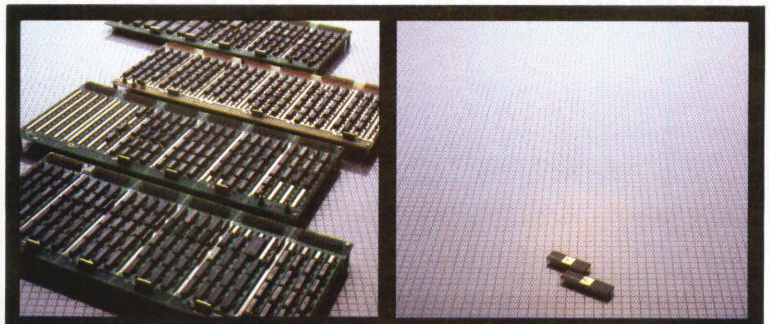
Joining forces for the ultimate system.

Mindset took advantage of our complete design environment with the best computer equipment and the finest IC software design tools available. Plus our top design team to work in tandem with their system designers.



It took teamwork to build the pyramids and teamwork to produce the Mindset – VTI's and Mindset's engineers worked in tandem using VTI's advanced IC design tools.

The resulting effort produced the heart of the Mindset system – two VLSI graphics co-processor chips. One dedicated to generating the bit-mapped graphics display. The other for performing bus arbitration between the graphics processor and the system processor. Together, they produce graphics operations at speeds of up to 50 times faster than a standard IBM PC.



Over 450 ICs from these four boards were squeezed onto two VLSI graphics co-processor chips – in only six months.

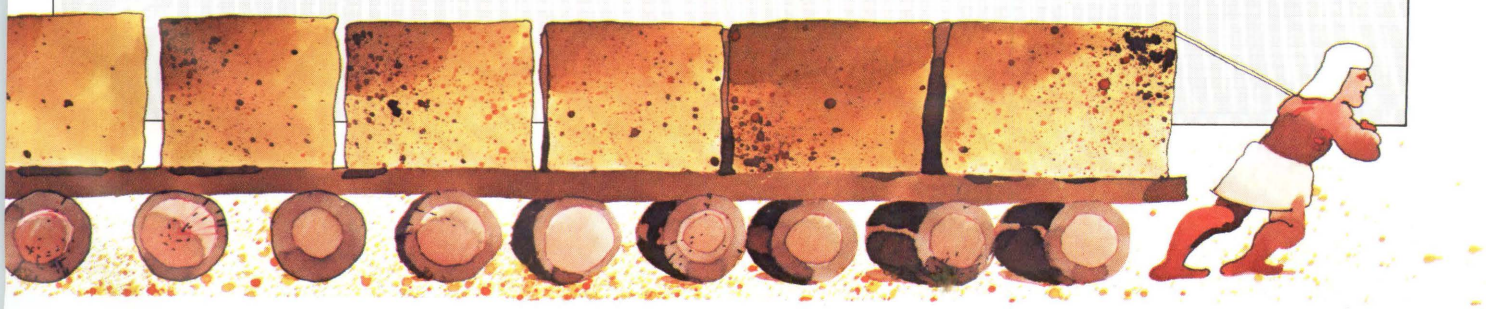
From 4 boards to 2 chips in just 6 months.

Considering the enormity of the task, turnaround time on the project was simply phenomenal. Working in conjunction with Mindset, the VTI team put 45,000 transistors onto a graphics chip and 25,000 transistors onto a bus controller chip. The entire design was accomplished in a record 19 weeks and the finished prototype in just six months.

System and chip design actually overlapped, enabling simultaneous enhancements in Mindset's architecture.

The net result – $50,000 graphics in an $1,800 package.

The bottom line was the chips worked. If Mindset had decided to use a standard approach, they wouldn't be one of the most successful new PC start-ups in the industry. Instead, they can offer engineering and business users a unique capability at a remarkable price. Thanks to some highly innovative assistance from VTI.



FOLLOW THE LEADER AND YOU'LL END UP IN FRONT.

If you have an impossible custom design task facing you, VTI can make it happen. Besides Mindset, we've done it for others. Companies like Drivetec, Tandy, Granger Associates, and Wang Labs.



VTI's extremely accurate photolithography equipment is capable of processing sub-2-micron geometries.

Our cell-based custom approach is only one of our wide range of proven ASIC design solutions. For example, our recently introduced gate array and programmable logic products offer the lowest cost VLSI solutions to simpler design problems.

Design in your own neighborhood.

You needn't travel to the factory to take advantage of our IC design system. Our Design Centers throughout the U.S. and Europe provide training, tools and expert engineering consulting support.

We're memory makers, too.

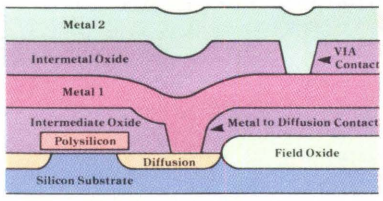
The VTI line of high performance Read Only Memories ranges from 32K to 256K. Advanced memory technology drives the processes on all of our products. So you're always assured state-of-the-art performance.

CMOS: The process of tomorrow is here today.

Our foundry service, the Silicon Express,™ accommodates your fabrication needs with high performance, 2-micron, double-metal CMOS. The process is used in our custom designs, gate arrays and many of our memory devices, too. With VTI you'll get the latest technology in your next generation product.

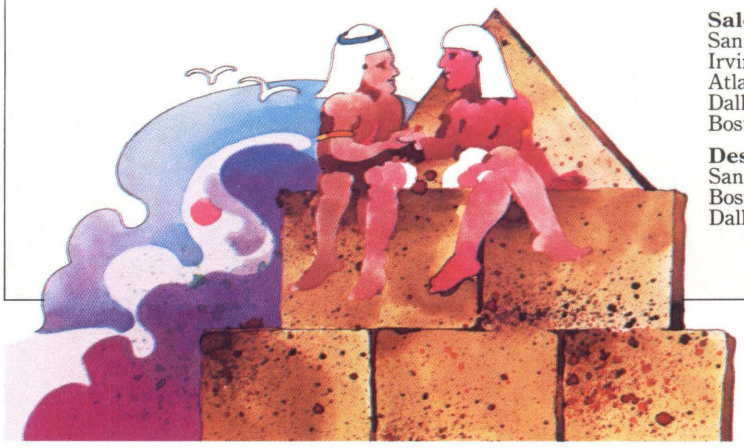
VTI really delivers silicon.

Last year, we produced over 14 million die. We'll double that this year, making us the exception to the industry rule of "sorry, no silicon." And our track record for quality exceeds industry standards. As you can see, VTI delivers on all counts.



2-micron double-metal CMOS technology: newly added to VTI's wafer processing capability.

Just as the pyramid builders of old were faced with an immense design problem, so are companies like Mindset. If you have a design problem in front of you, we can put it behind you. Fast. Contact any of our locations listed below, or call (408) 942-1810. Or, for more information, write VTI, Advertising Department, 1109 McKay Drive, San Jose, CA 95131.



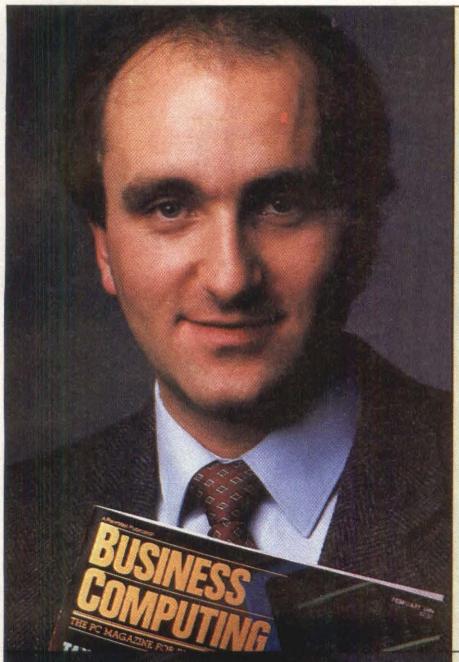
Sales Offices
San Jose, CA 408-942-1810
Irvine, CA 714-833-3103
Atlanta, GA 404-446-1326
Dallas, TX 214-231-6716
Boston, MA 617-229-6555

Design Centers
San Jose, CA 408-943-0264
Boston, MA 617-229-6555
Dallas, TX 214-231-6716

VLSI TECHNOLOGY, INC.

CIRCLE 100

"We couldn't manage without Business Computing."

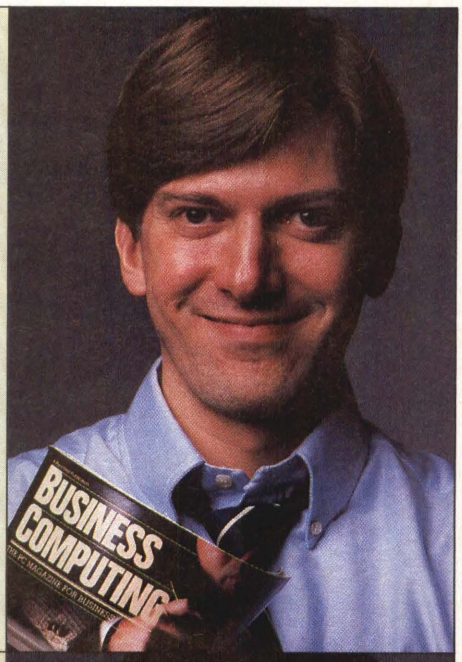


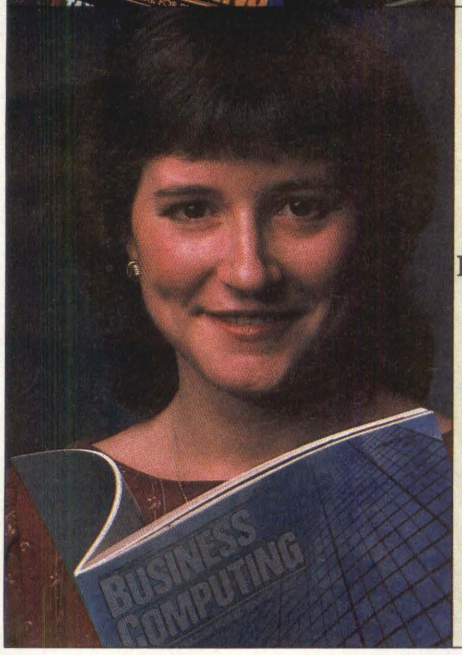
"Outlined vendor options for expanding our system."

Rolf Grueniger
PRESIDENT AND GENERAL MANAGER
Skytruck International Airfreight, Inc.

"Great overview for a manager trying to keep up."

Bill Fowkes
MARKETING MANAGER
Home Box Office, Inc.



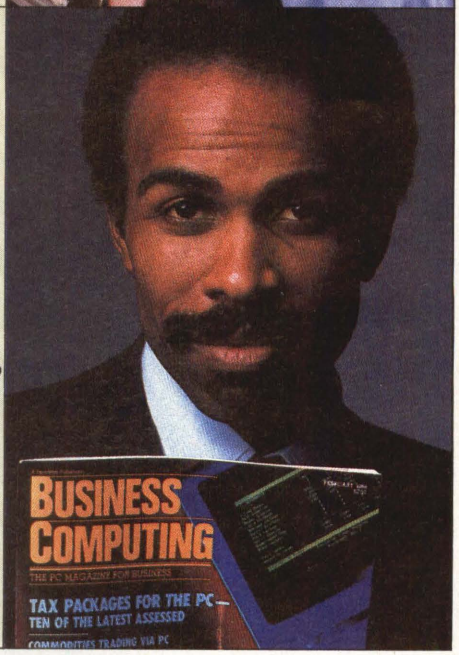


"Saved me hours of research on spread-sheet software."

Karen Leahy
PRODUCTION MANAGEMENT ASSOCIATE
Buckley and Urbanski

"Thanks for the sound advice on word processing."

Jim Hill
MANAGING EDITOR
Hill Publishing



Now there's a new magazine dedicated to helping you manage more profitably with your IBM PC. No games. No jargon. No fooling.

Instead, BUSINESS COMPUTING delivers executive-level wisdom every month on how to use the latest software and peripheral technology to make sharper decisions and solve management problems.

With step-by-step tutorials, sample applications and pertinent case histories. As well as guidance from managers like you, who share their own trial and error experiences so you won't have to.

BUSINESS COMPUTING. It's the first non-technical magazine for managers about using PCs for profit.

Try BUSINESS COMPUTING now. As a Charter Subscriber, 12 issues cost you only $14.75, over $9.00 off the regular subscription price. And BUSINESS COMPUTING is guaranteed to help you or your money back. Just use the reply card attached or call toll-free 1-800/922-4800. (In Utah, call 1-800-662-2500.)

BUSINESS COMPUTING

For the PC user who isn't playing games.

A Pennwell Publication/119 Russell Street, Littleton, MA. 01460

To order, circle Reader Service No. 139

This Publication is available in Microform.



Please send additional information for _____
(name of publication)

Name _____

Institution _____

Street _____

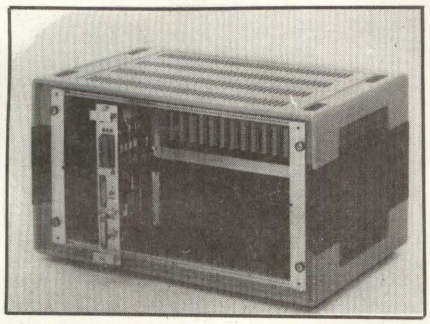
City _____

State _____ Zip _____

University Microfilms International

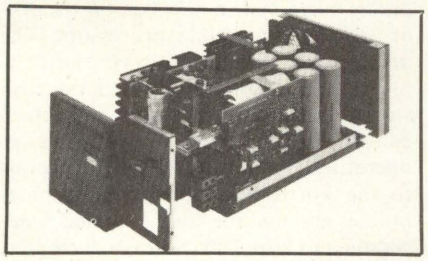
300 North Zeeb Road, Dept. P.R., Ann Arbor, Mi. 48106

Tabletop and rackmount chassis accommodates 20 VME boards



Including cabinet, rack, VME motherboard, power supply, and fans, the DSSERCK20 VMEbus chassis accepts up to 20 double Eurosized boards. This chassis is available in tabletop and 19-in. rackmount. At 5 V with 1 percent regulation, the DSSERCK20 supplies 45 A. Motherboard transfer rates reach 20 MHz. A companion unit, the DSSERCK09, accommodates up to 9 VME boards and can supply 20 A. Quantity prices start at $1995. **Data-Sud Systems/U.S., Inc**, 2219 S 48th St, Suite J, Tempe, AZ 85282.
Circle 388

Compact power supply handles remote drive, high volume needs



Extra output for burn-in where line load voltage drops are critical is provided by the SRX 5-200. This 5.5-V, 200-A power supply for offline switching is housed in an industry standard case of 5 x 8 x 11 in. (12.7 x 20.3 x 27.9 cm) and offers power density of 2.47 W/in.^{3} with output adjustable from 4.47 to 5.5 Vdc at 200 A maximum. Regulation is 0.1 percent, line and load. Ripple and noise measures 15 mV maximum rms. The SRX 5-200 weighs 15 lb (6.8 kg). Optional features include LED indicators, digital interfacing, reverse air flow, and external crowbar drive. Single units list at $850 with quantity discounts available. **Sorenson Co**, 676 Island Pond Rd, Manchester, NH 03103.
Circle 336

Packaging bridges gap between wire-wrap and multilayer boards

Unilayer II allows designers to go from prototyping to production, economically, in as little as two weeks. One Unilayer II board can equal the packaging density of any multilayer design so it can replace multilayer boards in final production. Software provides a direct conversion from wire-wrap without the risk involved in converting input data from one program to another. Typical cost of a wired, test socket board at the 25-piece quantity is approximately $3 per IC position. **Augat Inc, Interconnection Systems Div**, 40 Perry Ave, PO Box 1037, Attleboro, MA 02703. **Circle 337**

Brushless dc generators at home in harsh settings

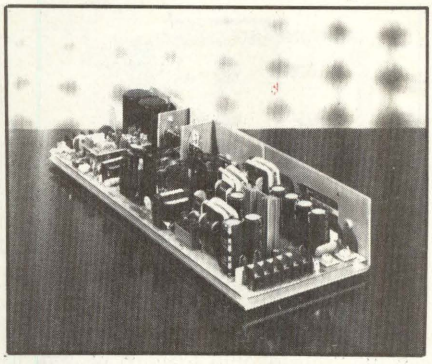
Brushless dc generators in the BDC series survive tough environments. In sizes from 5 to 250 kW, BDC series alternators boast resistance to Class 1, Group D, Div 2 environments. Required maintenance of the brushes and commutator is eliminated. Designs can be sized to a specific load; available voltages are 32, 125, 150, or 250 V. The standard BDC machine comes equipped with a voltage regulator that allows 10 and −30 percent voltage adjustment, and ±1 percent automatic voltage regulation, **Lima Energy Products**, PO Box 918, Lima, OH 45802.
Circle 338

Power MOSFET line delivers reduced on-state resistance

A line of high current TMOS power MOSFETs, which includes the MTE200N06, features low on-state resistance. Designed for high speed power switching applications (such as mainframe power supplies), these devices range from 50 to 200 V and exhibit 120 to 200 A continuous drain current ratings. MTE200N06 maximum on-state resistance hits 9 mΩ. For the MTE200N06, quantity 100 price is $120. **Motorola Semiconductor Products Inc**, PO Box 20912, Phoenix, AZ 85036.
Circle 339

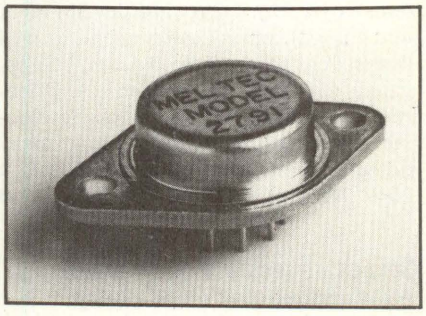
Let's hear from you
We welcome your comments about this issue. Just jot them on the Reader Inquiry Card.

Quad-output, open-frame switcher series offers selectable ac range



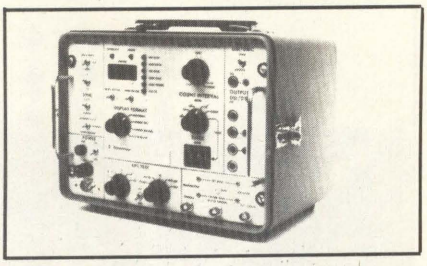
Quad-output switching power supplies in the MRX-400 series provide regulated 12- or 24-V, 1.5-A output for driving CRT displays. They also can supply 5-V, 50-A output as well as a choice of 12- or 15-V, 10-A outputs. Frames measure 2.5 x 5 x 15 in. (6.35 x 12.7 x 38.1 cm); soft-start circuitry minimizes input current surge at power-up. Selectable ac input may vary from 90 to 132 V or 180 to 264 V, 47 to 63 Hz. An adjustable main control loop regulates the line and load of the 5-V output to ±1 percent. Each supply weighs only 5 lb (2.268 kg). MRX-400 series supplies cost $407 each in purchases of 100. **Todd Products Corp**, 50 Emjay Blvd, Brentwood, NY 11717.
Circle 340

Wideband op amp uses monolithic JFET technology



Designed as a pin-compatible replacement for the μA791, the model 2791 wideband power operational amplifier incorporates internally compensated, matched junction-FETs. This provides low input bias (typically 30 pA) and offset current (3 pA). Typical gain bandwidth product is 5 MHz. The model 2791 costs $32 in quantities of 1000, and $65 for quantities of one to nine. **Mel Tec**, 411 Providence Hwy, Westwood, MA 02090.
Circle 341

Carrier analysis test set has self-contained receiver/transmitter



The model 273A portable error rate test set has a self-contained receiver and transmitter that offers complete testing for digital transmission systems as well as checking of T1 and T1C carrier, T1 outstate, and other systems. The 273A can provide checkout of 1.544 (± 50 bits/s in T1 mode) and 3.152 Mbits/s (± 50 bits/s in T1C mode), measure bipolar violations or logical errors, and inject controlled bipolar violations for testing of automatic protection switches. Sporting self-test capability, it can operate synchronously with an external clock. The device is priced at $2750. **Bowmar/Ali, Inc**, 351 Main St, PO Box 10, Acton, MA 01720.
Circle 342

Seven CAD and graphics systems expand CAE capabilities

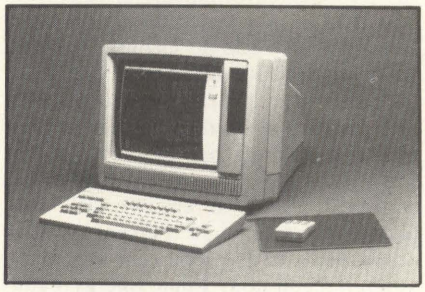
Seven interrelated graphics and CAD systems provide wide ranging CAD and CAE power. Configurable from a common Metheus 32-bit workstation base, the systems cover Unix software development and color graphics development, as well as design of schematics, logic, PC boards, gate arrays, and full custom ICs. Prices range from $24,900 for the Unix development system to $99,900 for the gate array design system. **Metheus Corp**, PO Box 1049, Hillsboro, OR 97123. **Circle 343**

Stripchart recorder has flexible input ranges and chart speeds

Multifunction keyboard controls and remotely activated operations mark the R100A stripchart recorder. This unit has 10 input voltage ranges and 12 chart speeds. A remote pulse generator can initiate its stepper motor chart drive, or it can operate in a timed mode. Other features include digital pen movement, automatic return and advance, and positive paper lock. **Perkin-Elmer Corp, Oak Brook Instrument Dept**, 2000 York Rd, Oak Brook, IL 60521. **Circle 344**

Workstation for CAE, based on 32-bit CPU, has 1-Mbyte RAM

The CDX-9000 CAE workstation offers documentation processing, hierarchical schematic design, and full function logic simulation. The workstation features a 32-bit CPU, 1 Mbyte of RAM, an 814-Kbyte (formatted) 5¼-in. floppy disk drive, a 35-Mbyte (formatted) Winchester disk, a 17-in. black and white monitor (1024 x 800), keyboard, and mouse. Berkeley 4.2 Unix with virtual memory is supported. A high end model—the Logic Design and Verification System—has 16 transistor models, over 95 logic primitives and an extensive library of TTL, CMOS, and ECL devices. The CDX-9000 is priced at $36,500 (quantity four to nine). The Logic Design and Verification System, available in a color version, costs $55,000 (quantity four to nine). **Cadnetix Corp**, 5797 Central Ave, Boulder, CO 80301.

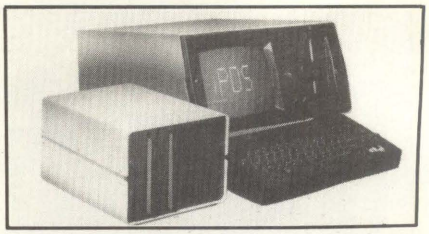


Circle 345

Graphics system extension packs high performance

Local hidden surface removal and smooth shading, plus transparent surfaces and automatic patterned dithering, upgrade the One/80, producing the model One/80-S. Dual-mode memory configuration is supported, which allows the One/80-S to be used as a high resolution 1280 x 1024 system or as a medium resolution 640 x 512 system for shaded image generation. Firmware automatically reconfigures image memory and performs dithering of the 24-bit color data to allow its display in an 8-bit deep image memory partition. Double buffering of shaded and Z-buffered images allows display of one 640 x 512 image while another is being constructed in an off-screen buffer. **Raster Technologies, Inc**, 9 Executive Park Dr, N Billerica, MA 01862.
Circle 346

Development system add-on ups memory capacity



Mass storage unit FDL-2 links to an Intel iPDS system to add 1.28 Mbytes of formatted data. The 8272 floppy disk controller chip is used to drive both the internal 5¼-in. drive and the FDL-2 add-on drives. These double-sided, double-density drives each hold 95 tracks/in. They configure as drive :F1: and :F2:, and are compatible with iPDS 100 integral drives. Flat-ribbon cable for installation is included. **IPD Electronic GmbH**, Im Steinbuhl 5, Industriegebiet Pinache, D-7135, Wiernsheim 2, Federal Republic of Germany.
Circle 347

Modular Forth language interpreter gives low software overhead

The MA2301 operates independently of disk drives and is designed to be used with the MA2000 family of Macrocomponents or any micro system based on the NSC8000- or 8080/Z80-compatible processors. The development system consists of the interpreter, resident compiler, RAM disk storage, 8080 assembler, and full screen editor. The system uses MVP-Forth as its operating system. MVP-Forth conforms to the Forth-79 standard and includes utilities and multitasking. **National Semiconductor Corp**, 2900 Semiconductor Dr, Santa Clara, CA 95051.
Circle 348

Software system integrates range of engineering applications

The computer integrated electronic engineering system allows users to create circuits on a workstation and simulate their performance. It can also support design of PC boards, hybrid, gate array, standard cell, and full custom circuits. VAX computers are used with V-series workstations for PC board design. Key feature is relational data base that enables data transfer between all workstations and software packages for consistent design and data. **Racal-Redac, Inc**, Lyberty Way, Westford, MA 01886.
Circle 349

System transforms PCs into design workstations

The Logicpro software package converts IBM PCs and XTs into engineering workstations for logic design. Designed for both board-level and IC logic designers, Logicpro converts the PC into a system that simulates all standard logic elements—for example, two libraries for CMOS and TTL each contain over a hundred 7400 and 4000 series models. The software can output all or partial network states at any time during simulation. A Logicpro/3000 version costs $650 and requires the 320-Kbyte system. TTLLIB and CMOSLIB cost $150 each. **E/Z CAD Inc**, 5589 Starcrest Dr, San Jose, CA 95123.
Circle 350

Plug-compatible system operates in standalone or distributed mode

Complete with its own 32-bit IBM-compatible computer and two high performance color graphics raster workstations, the System 9000 is intended for standalone or distributed CAD/CAM applications. It includes two display stations, a computer, 2 Mbytes of RAM, two 85-Mbyte disk drives, an IBM-compatible nine-track tape drive, a control terminal, and a printer. Each station is equipped with its own display and communications processor, keyboard, function keyboard, data tablet, and light pen. System lists for $120,000. **VG Systems**, 21300 Oxnard St, Woodland Hills, CA 91367.
Circle 351

Prototyping system speeds interface design

The FLAIR system promotes rapid prototyping of user interfaces. FLAIR (for Functional Language Articulated Interactive Resource) speeds development of user/machine interfaces, allowing designers to consider various formats, device choices, and sequences. The system sports compatibility with a variety of hardware I/O devices. Input methods include graphics tablet, light pen, joystick, and voice recognition. Supported output includes Ramtek, Megatek, and Tektronix products, plus VT 100 or compatible CRTs, and text-to-speech voice synthesizers. **TRW Inc**, One Space Pk, Redondo Beach, CA 90278. **Circle 352**

Hardware accelerator checks IC design rules

An integrated hardware option for the Graphics Design System II (GDSII), the Fast-Mask Engine improves, IC design rule verification. With 750 Kbytes of main memory and an 80-Mbyte Winchester disk, the Fast-Mask Engine off-loads the GDSII, while using advanced geometry checking and data compression. The Engine uses the same commands as the GPLII, which supports flexible checking of design cells, functional blocks, and complete circuit data. Prices begin at $60,000. **Calma Co**, 2901 Tasman Dr, Santa Clara, CA 95050. **Circle 353**

Four vital software packages



From one source . . . RTCS gives your PC/MS-DOS based systems professional development capabilities with a family of operating software tools. Call RTCS today for detailed information. We'll show you how you can extend the capabilities of your PC for less money than you'd expect.

RTCS/UDI —
UNIVERSAL DEVELOPMENT INTERFACE
Run Intel Series III software on the IBM PC/XT and other MS-DOS computers.
MS-DOS is a trademark of Microsoft

PC/iRMX
Real-Time, Multiuser, Multi-tasking Operating System for IBM PC and others.
iRMX is a trademark of Intel Corporation

INTEL LANGUAGES
Pascal 86/88 Compiler
Fortran 86/88 Compiler
PLM 86/88 Compiler
C 86/88 Compiler
ASM 86/88 Macro Assembler

OTHER UTILITIES
Software Debugger
Hardware Debugger
Target System Development Link

INTEL UTILITIES
Link 86 Linker
Loc 86 Locater

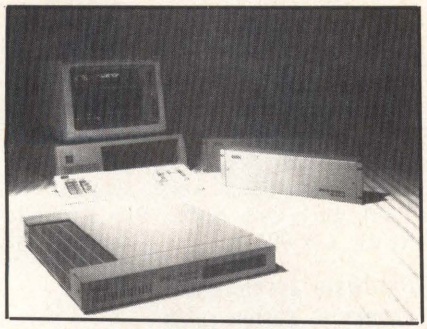
Distributed in: Japan by SYSCON CORP., Tokyo; Europe by MICRO-TECH ELECTRONICS, Paris; Israel by MLRN ELECTRONIC LTD., Zahia

RTCS REAL-TIME COMPUTER SCIENCE CORPORATION
SOFTWARE DEVELOPMENT TOOLS
ADVANCED OPERATING SYSTEMS

P.O. BOX 3000-886, CAMARILLO, CALIFORNIA 93011 • PHONE NO. (805) 987-9781 • TELEX 467897

CIRCLE 123

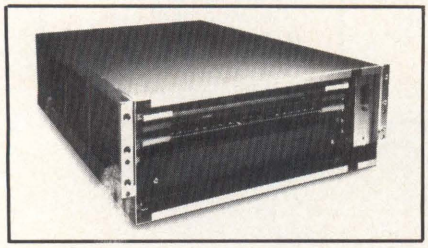
Controller lets micros perform data acquisition and control



Acting as a frontend processor interfacing varied analog and digital data to microcomputer buses, the PCI-3000 lets smaller machines perform work previously done by mainframes and minis. Compatible hosts include computers from Apple, Compaq, DEC, and Hewlett-Packard, as well as the IBM PC and XT. Users can access up to 192 I/O points from a single master unit. The PCI-3000 is capable of RS-232, RS-422, and IEEE 488 communications. Compact master unit measures 2½- x 17- x 19-in. (6.35- x 43.2- x 48.3-cm). Price is $3000. **Burr-Brown,** International Airport Industrial Park, PO Box 11400, Tucson, AZ 85734.
Circle 354

Rackmount terminals handle industrial color graphics

Designed for process control and industrial settings, the 6210 family of rackmount terminals offers a choice of three performance levels in a compact 7- x 19- x 21-in. (17.8- x 48.3- x 53.3-cm) package. The low end R6210/01 provides complete VT100 terminal emulation and meets ANSI 3.64 protocol standards. The high end R6210/21 has all R6210/01 terminal features, plus an overlay with four planes of bit-mapped graphics at 640 x 480 resolution. System can display up to 16 colors from a palette of 64. Prices for the system are R6210/01, $3495; R6210/11, $3995; and, R6210/21, $4995. **Ramtek Corp,** 2211 Lawson Lane, Santa Clara, CA 95050.



Circle 355

Workstation runs 3-D vision/modeling and robot programming

As a CAE workstation, Visicam generates 3-D vision programs for inspection and robotic applications. The vision programs run on the Silma 9000 vision processor. Another workstation, the Robocam, is designed for online or offline programming of industrial robots, and for simulation of robots and their workcells. Host hardware is a dedicated 32-bit desktop mainframe with high resolution graphics. A typical Visicam workstation, including an Apollo DN 300 workstation (70-Mbyte desk unit), an SDSP 80 vision processor, a CCD camera, a color monitor, and 3-D vision software, is $80,500. With both Robocam and Visicam, price is $107,500. **Silma Inc,** 1800 Embarcadero Rd, Palo Alto, CA 94303.
Circle 356

Integrated package aids development of numerical control system

A closed-loop interpolator/servo control package allows users to design a control system around the 8086/8088. The package includes the I/S-86 motion control operating system and the 9702 dual motion control interface modules. The user matches this with a Multibus processor and memory. The I/S-86 software provides up to 8 axes of simultaneous motion, linear, circular, and/or helical interpolation, and compatibility with Isis and iRMX languages. The interface module provides an independent quadrature encoder plus marker decoding, and analog servo amp drive output. **North Coast Automation Inc,** 71 Alpha Pk, Cleveland, OH 44143.
Circle 357

Bus controller supervises remote instruments

Housed in a desktop VDU unit incorporating twin 5¼-in. floppy disk drives and detachable keyboard, the Type SE2650 controls any IEEE 488 device. This bus controller can configure to drive single or multiple instruments. Program generation and data storage capacity hit 128 Kbytes of RAM, while the twin floppy drives provide 320 Kbytes of permanent storage. Fully configurable RS-232-C, Centronics, and IEEE interfaces are standard. The system supports Apple DOS 3.3. **Thorn EMI Datatech Ltd,** Spur Rd, Feltham, Middlesex TW140TD, England.
Circle 358

Processor takes care of robotic vision and visual inspection



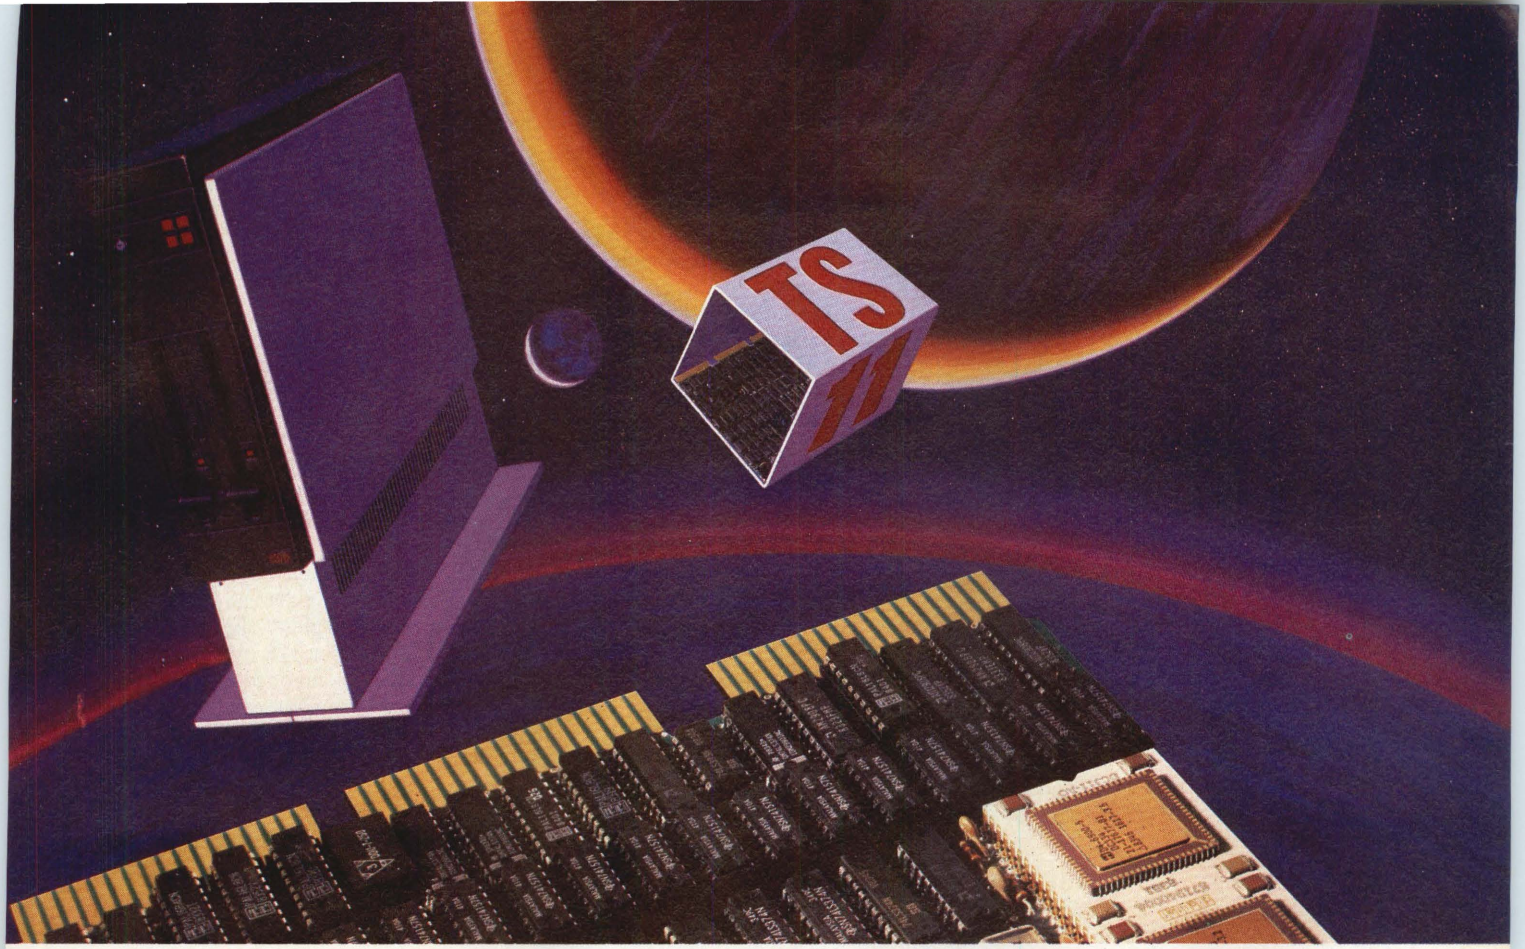
Odin 20 machine vision processors operate in real time and accurately execute dimensional measurement, pattern recognition, automated sorting, and online guidance and control. Working at 40 MIPS, the Odin 20 processes image data at a spatial resolution of up to 2048 x 2048 pixels; selectable radiometric resolution reaches 256 gray levels. The 1.5 ft^3 unit mounts on standard 19-in. rack and accepts multiformat sensor inputs. Basic configuration is $29,000, with multiple unit discounts available. **International Imaging Systems,** 1500 Buckeye Dr, Milpitas, CA 95035.
Circle 359

Micro-based system works in tough environments

Diverse optional memory boards, iPC Basic, and a development system-on-a-board, comprise the industrial personal computer system. A 64-Kbyte, battery-backed RAM/ROM board handles large storage operations, while the iPC-MBM bubble memory board is available for especially hostile environments. All development code in an iPC-System is permanently stored in ROM, where it is immune to most hazards. Where disk drives are needed for data collection, the iPC-31 microfloppy-on-a-board provides hard-jacket, shutter protected media. **Vesta Technology, Inc,** 2849 W 35th Ave, Denver, CO 80211. **Circle 360**

Industrial computer is PC compatible

The RM-1600 is housed in a rackmountable chassis so it can be mounted with other equipment in a standard modular 19-in. rack. A 9-in. video monitor is included in the system housing. Its built-in 640 x 325 resolution graphics have more than 50 percent greater resolution than the IBM PC. The unit also includes two floppy drives. The 1600-10 version is available with a 10-Mbyte hard disk and a floppy drive. **Materials Development Corp,** 21541 Nordhoff St, Chatsworth, CA 91311. **Circle 361**



Get On-board the J11 Breakthrough.

We didn't invent the J11—we just took advantage of the technological breakthrough it represents to create the most powerful front end communications processor for Q-bus* applications in the world. It's called the MLSI-JFEP11, and now, for the first time, co-processing on the Q-bus is possible. And with this multi-processing capability, comes a substantial cost savings for system designers as a new world of applications are available to the smaller, yet faster, CPU based systems.

Powerful is an understatement. The MLSI-JFEP11 is the first application of the DEC J11* processor other than on the 11/73 single board computer. It features 512KB dual ported memories, two high speed serial ports (one megabaud each) and an external parallel bus for I/O expansion.

And getting that power underway has never been easier. The MLSI-JFEP11 is programmed in the most widely used machine language in the world—the PDP-11* instruction set. Needless to say, that means a rapid transportation of countless existing programs and a minimum start-up time for new ones.

So if you're committed to the new micro architecture with macro capabilities, we should travel in the same circles. Especially when it comes to such system applications as dynamic communication line resource allocation, data compression/decompression, message processing, message routing, protocol conversion and multi-processing.

Because the MLSI-JFEP11 is part of the new world of MDB capabilities. One that started with the TS11 Controller/Coupler, followed by our MICRO/11 and MICRO/32 packaging breakthroughs. All tools to help you go as far in the building of a super micro computer system as your design imagination allows.

*Q-bus. DEC J11 and PDP-11 are Trademarks of Digital Equipment Corporation

MDB SYSTEMS INC.

Corporate Headquarters
1995 N. Batavia Street, Box 5508
Orange, CA 92667-0508
Tel. 714-998-6900 TWX: 910-593-1339
FAX: 714-637-4060

MDB Systems U.K., Ltd. Berkshire, Tel. 44 06286 67377
MDB Systems GmbH Munich, Tel. 49 89 9101272
MDB Systems Australia Neutral Bay, Tel. 612 929-6526

In the western states MDB products are also distributed by WYLE LABORATORIES ELECTRONICS MARKETING GROUP

See us at MINI MICRO SOUTHWEST and the Federal Computer Conference

CIRCLE 124

Slave processors exhibit 4- or 6-MHz operation

Capable of interfacing with any 696.1/D2 master with extended memory capacity, CPS-MX series slave processors feature options of 4 MHz or 6 MHz processor speeds and 64 or 128 Kbytes of bank-selectable memory. These processors are designed for use with TurboDOS 1.3 and CPZ-4800 master processors. The CPS-MX series provides two serial synchronous or asynchronous I/O ports, two parallel ports and software selectable baud rates. Discrete refresh circuitry releases the CPU from memory refresh responsibility. **Intercontinental Micro Systems**, 4015 Leaverton Ct, Anaheim, CA 92807.
Circle 362



... goodbye to backup

Now you can put big system disk storage capacity in your micro- or mini-based system. Up to 106 Mbytes. For less than $10,000*. And half that capacity is removable media. So you can say goodbye to all those hours of slow backup. And goodbye to the added cost of separate tape or disk backup drives.

HD-26 also reduces disk transfer overhead with its average access time of 35 msec.

The HD-26 is compatible with the industry standard SMD interface. It's CDC Lark® compatible. And a version is optionally available with an IEEE 696 (S-100 bus) controller. Including cables and host software if your system is Cromix based Cromemco.

HD-26 is also available in a 53 Mbyte configuration for less than $6000*.

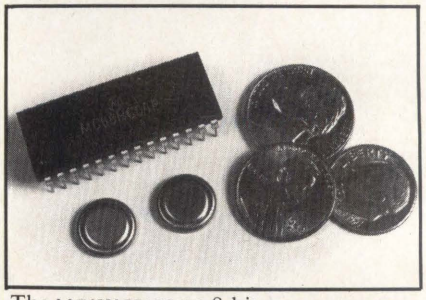
* Quantity price with SMD interface
® CDC Lark is a registered trademark of Control Data Corporation

For more information on our complete line of disk systems, call us at 703-281-4666.

VCA

VCA CORPORATION
501 Church Street, NE Vienna, VA 22180
(703) 281-4666 Telex. 469088

Compatible with TTL/CMOS, chips have dual I/O



The MC68HC04P2/P3 8-bit computers on a single chip contain CPU, onchip clock, 1- or 2-Kbyte ROM (plus 72 bytes for lookup tables), and 32 or 128 bytes of RAM, as well as memory mapped I/O. With software features similar to those of M6800 family members, these devices sport 20 TTL/CMOS-compatible bidirectional I/O lines (eight LED-compatible lines). Volume pricing is approximately $4. **Motorola Inc, MOS Microprocessor Div**, 3501 Ed Bluestein Blvd, Austin, TX 78721.
Circle 363

Board extends 32-bit processing for Multibus applications

The GVC-32 Multibus board aims at multiprocessing applications that require 32-bit power. Based on the NS32032, this board includes 0.5 Mbyte of dual-port RAM, allowing message passing by another processor, or DMA by a disk controller. A mailbox interrupt informs the GVC-32 when another processor posts a message in its dual-port RAM. A 10-MHz processor can execute without wait states, and without using Multibus bandwidth. Execution can occur over the Multibus as in conventional systems (in which case, the GVC-32 automatically performs two 16-bit Multibus read cycles for a single 32-bit fetch). Price is $3995. **GVC Microcomputers Inc**, 222 Third St, Cambridge, MA 02142. **Circle 364**

Coprocessor card delivers CP/M capabilities to the VAX

The D200 is an 8088-based coprocessor card for PDP-11 and VAX computers. It comes with 256 Kbytes of memory and is packaged with the CP/M-86 operating system. Using this combination, CP/M-80 or CP/M-86 programs can run on VAX and PDP-11 equipment, and peripherals can be shared. The D100, a Z80 processor card, is required to run the D200. The D100 costs between $1100 and $1895, depending on host system configuration. The D200 costs $1195. **Decmation**, 3375 Scott Blvd, Santa Clara, CA 95051. **Circle 365**

OEM Decision Makers

"Only the Invitational Computer Conferences bring the latest OEM computer and peripheral products to your front door. You'll find us there!"



KARLH KUEBERT, General Manager, Shugart GmbH. W. Germany

And you'll find other top OEM manufacturers, such as Bull Peripheriques, IBIS, Fujitsu, Tandberg, Tandon, Olivetti and Control Data, to name a few.

In their 4th year, the "OEM Only" Invitational Computer Conferences bring you, the volume buying decision makers, together with the key suppliers of computer and peripheral products. The ICCs, a series of six, one-day regional shows are convenient to where you live and work. The social business setting makes it easy for you to meet potential suppliers one-on-one, and attend high tech seminars of your choice. As an invited guest, there is no cost to you.

Hear what the OEM manufacturers have to say, learn more about their products, and remember, you may attend "by invitation only."

1984/85 Europe ICC Locations

Oct. 9, '84	Munich, W. Germany
Oct. 16, '84	Vienna, Austria
Oct. 23, '84	Milan, Italy
Feb. 28, '85	Paris, France
Mar. 7, '85	Frankfurt, W. Germany
Mar. 14, '85	London, England

Call your local OEM supplier for your invitation or fill out the coupon and mail to:

B. J. Johnson & Associates, Inc.
3151 Airway Ave., #C-2
Costa Mesa, CA 92626
Phone: (714) 957-0171
Telex: 188747 TAB IRIN

Yes! I need an invitation to your "OEM Only" ICC. The nearest ICC to me is: _____

I buy in volume:

- ☐ Computers
- ☐ Disk/Tape Drives
- ☐ Controllers/Interfaces
- ☐ Terminals/Graphic Displays
- ☐ Software
- ☐ Printers
- ☐ Memory Boards
- ☐ Modems/Multiplexers
- ☐ Power Supplies

Name _____

Title _____

Company/Division _____

Address _____

City _____ State _____ Zip _____

Mail To: B.J. Johnson & Associates, Inc., 3151 Airway Avenue, #C-2, Costa Mesa, CA 92626
Phone: (714) 957-0171 Telex: 188747 TAB IRIN

CIRCLE 126

Speech synthesizer has internal controller

Voice-Chip speech synthesizer requires no external controller or address encoder since it carries those functions onchip. This 40-pin chip uses the DataVoice encoding process that assures up to 16 s of speech. It can be matched with memories of 32, 64, or 128 Kbits. Two-week turnaround for custom words and phrases is available. The Voice-Chip costs $7 in quantity. **DataVoice Corp,** 2 N LaSalle, Suite 1900, Chicago, IL 60602.
Circle 366

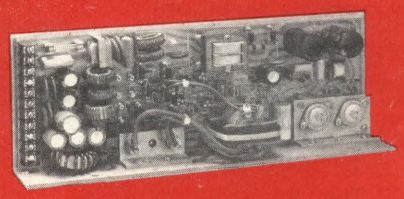
POWER SUPPLIES

LINEARS - SWITCHERS

LOWEST PRICES in USA!

359 Models - 1 to 4 Outputs - Open and Closed Frames - Form, Fit and Function for:

Boschert
Powertec
Power-One
Power/Mate
Lambda
ACDC
Elpac
Condor
LH



If You're Not Using Deltron – You're Paying Too Much!

Call Toll Free 800-523-2332

AC 43

Deltron inc. P.O. BOX 1369 WISSAHICKON AVENUE ■ NORTH WALES, PA 19454 PHONE: 215/699-9261 TWX: 510/661-8061

Standalone chip offers error correction and fast cycle times

The SN74ALS632 32-bit parallel error detection and correction circuit drives buses directly without the need for external bus drivers. The chip detects and corrects single bit errors, and detects and flags double-bit errors. They also flag gross error conditions of all high or all low outputs. The circuit can indicate a single-bit error in 40 ns and a double-bit error in 60-ns maximum. Corrected output data appears on the bus in 58 ns after processing. The chip provides byte writing capability and three-state outputs. Price in 1000s is $79.42. **Texas Instruments, Inc, Semiconductor Group,** PO Box 809066, Dallas, TX 75240.
Circle 367

Protocol engine uses shared memory to drive communications

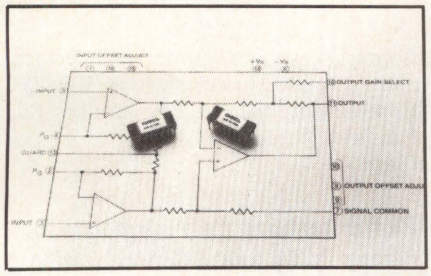
Family members in the QM10 monolithic communication controller group implement complex, upper-level protocols in firmware. The QM10 is a general-purpose protocol engine that applies the Transmission Control Protocol/Internet Protocol and the Internet Control Message Protocol. Communication takes place via shared memory and hardware strobes. This 40-pin device has a ROM that connects in piggyback fashion. A shared, multiport memory divides into two main address spaces (user and network) each totaling 8192 bytes. Maximum input voltage measures −0.3 to −7 Vdc. **Communication Machinery Corp,** 1421 State St, Santa Barbara, CA 93101.
Circle 368

Fast single-chip PROM shortens cycles and decreases power use

Sporting a clock-to-output time of 15 ns, the 63RA481A PROM features synchronous and asynchronous 3-state enable inputs plus asynchronous preset and clear. A related 63RA481 version of this chip has clock-to-output time of 20 ns. Both these 4-Kbit memories have onchip D-type registers. Apt for applications such as microprogram control storage, state sequencers, next address generation, and bit mapping, the chips draw a typical supply current of 130 mA and 180 mA, worst case. Power dissipation is 900 mW. Organization is 512 words x 8 bits. In quantity 100, the 63RA481 costs $11.20 and the 63RA481A, $15.40. **Monolithic Memories,** 2175 Mission College Blvd, Santa Clara, CA 95050. **Circle 369**

CIRCLE 127

Amplifier meets applications requiring fast setting



Programmable gain instrumentation amplifiers in the AM-551 line possess user-selectable gains of 1 to 1000. Settling time is 2 μs for a 20-V step to 0.01 percent accuracy and an offset voltage drift maximum of ± 15 μV/°C. Slew rate measures 23 V/μs with a small signal bandwidth of 400 kHz. These functionally complete devices consist of a high impedance, variable gain, voltage follower input stage, and a differential output stage. Operation over a wide range of temperatures is available on diverse AM-551 models. **GE-Datel,** 11 Cabot Blvd, Mansfield, MA 02048.
Circle 370

Filter/frequency detectors form telecommunications chip set

The S3526M bandpass/notch filter merges with the 8-pin S3524A digital frequency detector to form a 2600 Hz signal detector. This CMOS chip set manages telecommunication signaling, though individual chips can be used for filtering and frequency detection. Typically, these clock tunable devices cover 100 Hz to 5 kHz and require 100 W. The 19-pin S3526M is a tunable sharp band-pass and notch (band-reject) filter with an onchip tone generator. The chip set's notch cutoff frequencies permit a wide range of speakers' voices. S3526M price is $13.20 per 100 devices. The S3524A costs $5.50 per 100. **American Microsystems, Inc,** 3800 Homestead Rd, Santa Clara, CA 95051.
Circle 371

Multiplier combines high speed and low power

The MPY-16IMI high speed CMOS 16- x 16-bit parallel multiplier operates at speeds comparable to similar bipolar devices, but at 1/20th the power. It features a typical clocked multiply time of no more than 170 ns and operates with a power consumption of less than 150 mW at 5 V. Supply voltage ranges from 3 to 12 V. In quantities of 1 to 10, the chips are available at $69 each. **International Microcircuits Inc,** 3350 Scott Blvd, Bldg 37, Santa Clara, CA 95051.
Circle 372

Circuit tackles complex logic functions for design flexibility

A 40-pin PAL (PAL32R16) has 32 array inputs and 16 outputs. The 1500-gate equivalent device features product term sharing, programmable bypass, and an output polarity option. Preload and auto-test vector generation achieve complete testability. The product term sharing feature allows 16 product terms to be shared between two outputs; the bypass feature lets output registers be bypassed in bands of eight, leaving combinatorial outputs. Price is $53.94 in quantities of 100. **Monolithic Memories,** 2175 Mission College Blvd, Santa Clara, CA 95050.
Circle 373

PERIPHERALS

High speed nonimpacting printer reduces printing costs

The Anser 1 printing system prints data and text of almost unrestricted size, shape, and orientation directly from digital information at speeds to 125 pages/min. The system has a high speed, onboard, bit-mapped computer to allow users to control literally every dot on a maximum page size of 9 x 14 in. It can be connected to a host computer remotely controlled by a modem, or be driven by an integrated tape drive in offline mode. Price is $149,000. **Anser Technology, Inc,** 5535 Airport Freeway, Fort Worth, TX 76117.
Circle 374

Touch-screen kit attaches to IBM PC color monitor

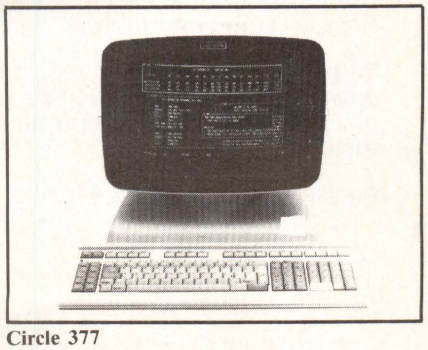
The Point-1 color kit includes a 13-in. diagonal touch screen, intelligent controller, and an RS-232-C serial interface. Fully programmable, the kit is supported by a complete set of software development tools, including a color release of the MicroTouch View*Point software. With the screen, users can select from menus, position the cursor, and create and manipulate graphics by touching the monitor with a fingertip. It offers a 1024 x 1024 touchpoint resolution. In quantity, pricing is under $650. **MicroTouch Systems, Inc,** 400 W Cummings Pk, Woburn, MA 01801.
Circle 375

Ergonomic display terminals boast 1000 different chars in one unit

The TDV 2200 S is an addition to a line of smart/editing terminals. Enhancements include a 70-Hz refresh rate for a completely flicker-free and stable picture and positive and negative graphics rendition. In addition, the terminals can connect to most mainframes including DEC, IBM, Honeywell, and Norsk Data. The unit can hold up to 56 Kbytes of memory. It can send or receive data in asynchronous, synchronous, or HDLC mode via RS-232 or RS-422 interfaces. **Tandberg Data Inc,** PO Box 99, Labriola Ct, Armonk, NY 10504.
Circle 376

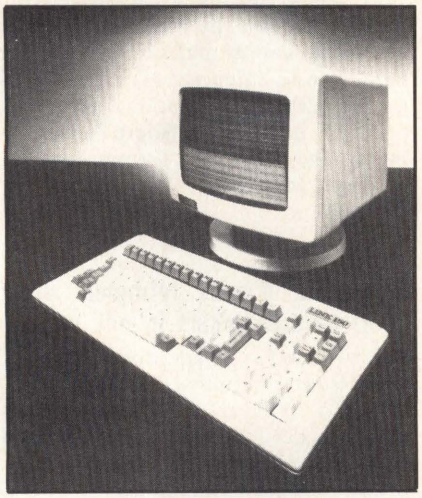
Computer display terminal features four-session windowing

The model 1221 Open Window allows a display station user to view four windows at once with each window representing a fully interactive session on a mainframe, mini, or other host. The unit has a keystroke record-playback feature, horizontal and vertical scrolling, and a modifiable keyboard. Using a series 400 controller, the unit can be connected to one or two IBM mainframes or up to 16 non-IBM systems. Screen sizes include 24 x 80, 32 x 80, 43 x 80, or 27 x 132 (all lines x columns). Prices range from $2066 to $2866. **Lee Data Corp,** 7075 Flying Cloud Dr, Minneapolis, MN 55344.



Circle 377

Video display terminal highlights operator convenience



Compatible with the TeleVideo 950, the Smart Link 150 display terminal offers ergonomic design highlighted by operator convenience features that include tilt-and-swivel display, 12-in. nonglare screen, and detached, low profile DIN standard keyboard. The terminal is optionally available with a special WordStar package. Depressing one key converts the terminal's program to work with systems running WordStar. The Link 150 is available with four pages of memory and 11 user-programmable function keys with nonvolatile memory. Quantity one: $895. **Link Technologies Inc**, 1887 O'Toole Ave, San Jose, CA 95131. **Circle 378**

Optical character recognition system reads and transmits to host

The FormsReader speeds data entry on mainframes made by IBM, DEC, Wang, Honeywell, NCR, HP, Data General, and others. The system selects, reads, validates, and edits typed, numeric, hand-printed, or handmarked information from preprinted forms. It is then transmitted to a computer, 3741-compatible floppy disk, or 9-track magnetic tape for processing. The system consists of a page reader, a micro with a 24-line CRT, dual floppies, keyboard, printer, an interface, and software. Prices start at $39,400. **CompuScan, Inc**, 81 Two Bridges Rd, Fairfield, NJ 07006.
Circle 379

Plotter features DM/PL software compatibility and 4-G acceleration

The DMP-51 plotter combines a high speed of 22 in./s, 4-G acceleration, and a resolution of 0.01 in., while producing C- or D-size drawings. With 26 Kbytes of built-in intelligence, the DMP-51 executes complex graphics functions from simple commands. It is compatible with existing DM/PL programs and uses servo-motor technology. A mechanical/architectural version, the DMP-52 [with 18- x 24-in. (45.7- x 61.0-cm) and 24- x 36-in. (61.0- x 91.4-cm) paper sizes], is also available. Both plotters are priced at $4495. **Houston Instrument**, PO Box 15720, Austin, TX 78761.
Circle 380

NOTICE

Request For Information
Data Encryption Technology

The Federal Reserve System is preparing a Request for Information (RFI) about the availability of encryption hardware and software. The Federal Reserve operates a large, national data communications network that is compatible with IBM's System Network Architecture (SNA). Encryption technology manufacturers and vendors may request to receive the RFI package, which will be distributed on or about August 15th, 1984. This is not a solicitation; it is a notice only.

All requests to receive the RFI must be in writing; telephone calls cannot be accepted. Requests must be accompanied by a copy of the latest Annual Report or a recent financial statement of the firm making the inquiry.

Requests must be received at the Federal Reserve Bank of Dallas within 15 calendar days following publication of this notice. A direct reference to the publication and its date of issue must be included.

Please send your request to: Automation Program Office
Federal Reserve Bank of Dallas
400 South Akard Street
Dallas, Texas 75222

Attention: Encryption Project

CIRCLE 128



TAUBER AND GATES, YOUR POWER CONNECTION.

Gates energy cells give you design flexibility, long life, superior performance and easy charging. Basic cells are built up in any configuration; series or parallel, case or shrink wrap. Batteries are designed to meet the needs of your application, not the other way around.

Tauber Electronics helped pioneer the Gates Value Distributor concept. From prototype to O.E.M. production, we deliver.

Gates
Energy Products

TAUBER ELECTRONICS INC
4901 Morena Blvd. ste. 314
San Diego, CA 92117
619/274-7242

CIRCLE 129

Compact printer aimed at portable applications

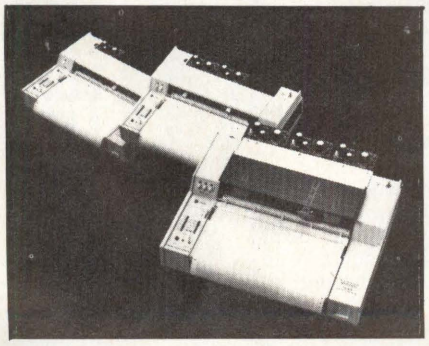


Measuring 13 x 7.5 x 2.8 in. (33 x 19 x 7.1 cm) and weighing 6.6 lb (2.9 kg), the GLP printer is intended as a portable computer companion. It features IBM PC compatibility and PC block graphics, as well as enlarged, condensed, emphasized, and double-strike print modes. The GLP is available with a Centronics parallel interface, or with both RS-232 and Centronics interfaces combined in one unit. Price is $299. **Centronics**, 1 Wall St, Hudson, NH 03051.

Circle 381

Lab recorders use single-, dual-, and three-pen format

The Soltec 1240 flatbed recorder comes in single-, dual-, and three-pen configurations. It features 23 chart speeds and 17 input ranges (from 1 mV to 200 V). The 1240 series has plug-in input spans, which allow optional temperature and current measurement. Event marker is superimposed on the trace of this 250-mm chart width unit. **Soltec Corp**, 11684 Pendleton St, Sun Valley, CA 91352.



Circle 382

Plotter option allows automatic changing of various pens

The Alphaplot II, equipped with a pen changing option, is a large format plotter. The pen-changing mechanism handles up to six HP-compatible plotter pens. An automatic pen-capping mechanism allows wet ink technical drawing pens, fiber tip pens, and ball tip pens to be changed automatically. The pen changer can be simply added to existing Alphaplot IIs as a field upgrade. Alphaplot II price, with pen changer, is $5990. Pen changer upgrade is $495. **Alpha Merics Corp**, 20931 Nordhoff St, Chatsworth, CA 91311.

Circle 383

Graphics and alphanumeric capabilities mark 3274-like terminals

Cluster controllers and terminals in the ID-200 series work like IBM 3270, 3271, and 3274 cluster controllers. Connecting via an RS-232-C interface, the units offer advanced graphics and alphanumeric traits, plus asynchronous, bisynchronous, and SNA/SDLC communication capabilities. Display list feature allows storage of often used file and command combinations in the terminal's display list memory for quick recall. Ergonomic design allows the monitor to tilt 15° forward and backward, and swivel 60° left or right. The ID-200 has 1280 x 480 resolution, scaling up to 32-K x 16-K pan, 16 levels of zoom, and 8 or 16 colors. **ID Systems Corp**, 4089 Leap Rd, Hilliard, OH 43026.

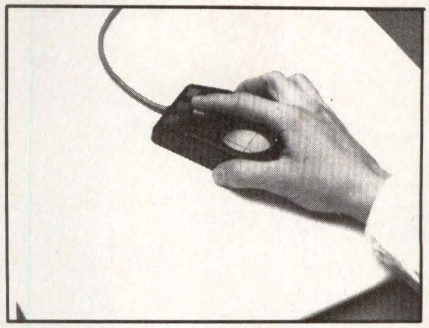
Circle 384

Second-generation imager captures computer graphics

Enabling users to make instant color prints of images from a 9-, 12-, 13-, or 19-in. CRT, the Instagraphic CRT imager is a low cost modular system. The device offers a variety of cone adapters to match almost any screen size. A print module includes the shutter and optical elements with a variable focus lens. It provides the correct lens-to-screen focal length and partially corrects for screen distortion. An adapter bracket allows the user to replace the print module and camera back with a 35-mm SLR camera. Cost is less than $300; individual cone adapters are less than $40 each. **Eastman Kodak Co**, 343 State St, Rochester, NY 14650.

Circle 385

Mouse-size cursor improves CAD station digitizing

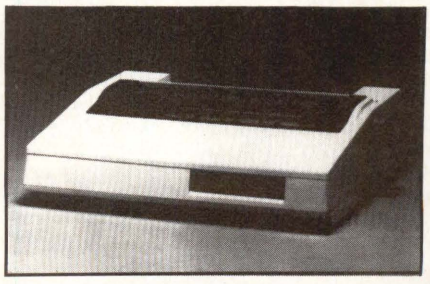


Tracing, menuing, and pointing are streamlined by the DP5-4CX and MD7-4CX cursors. The former is compatible with the Digi-Pad electromagnetic digitizer, the latter with the Micro Digi-Pad. Fitting comfortably in the user's hand, these ergonomically designed devices sport precision cross hair reticles. The cursors have color-coded keys arranged for compatibility with one- and four-button digitizing cursors and a three-button mouse. In 100s, prices are $210 for the DP5-4CX and $120 for the MD7-4CX. **GTCO Corp**, 1055 First St, Rockville, MD, 20850.

Circle 386

Dot-matrix units are IBM PC compatible

The low end 7500E is an IBM compatible, 80-col printer with 105-char/s speed and 45-lines/min output; its high end mate, the 8510/1550SCE, is a PC-compatible version of the two-speed, seven-color 8510/1550SC printers. The 7500E has a 9 x 9 matrix for dot-addressable graphics, plus friction and tractor paper feed, and bidirectional printing. The 8510/1550SCE prints graphics at 120 chars/s, and rough drafts at 180 chars/s. Throughput on these models hits 100 lines/min on the 8510SCE and over 45 lines/min on the 1550SCE. Price is $450 for the 7500E, $950 for the 8510SCE, and $1270 for the 1550SCE. **C. Itoh Electronics, Inc**, 5201 Beethoven St, Los Angeles, CA 90066.



Circle 387



Photo: Peter B. Kaplan

If you still believe in me, save me.

For nearly a hundred years, the Statue of Liberty has been America's most powerful symbol of freedom and hope. Today the corrosive action of almost a century of weather and salt air has eaten away at the iron framework; etched holes in the copper exterior.

On Ellis Island, where the ancestors of nearly half of all Americans first stepped onto American soil, the Immigration Center is now a hollow ruin.

Inspiring plans have been developed to restore the Statue and to create on Ellis Island a permanent museum celebrating the ethnic diversity of this country of immigrants. But unless restoration is begun now, these two landmarks in our nation's heritage could be closed at the very time America is celebrating their hundredth anniversaries. The 230 million dollars needed to carry out the work is needed now.

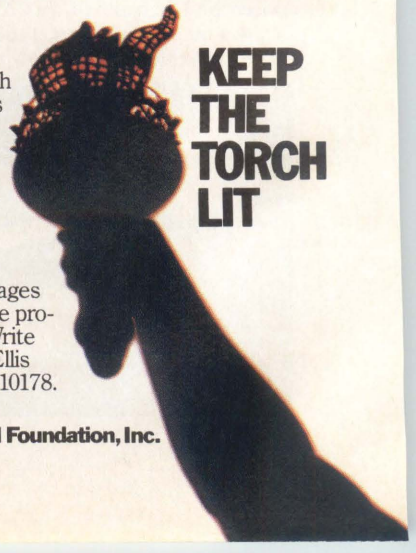
All of the money must come from private donations; the federal government is not raising the funds. This is consistent with the Statue's origins. The French people paid for its creation themselves. And America's businesses spearheaded the public contributions that were needed for its construction and for the pedestal.

The torch of liberty is everyone's to cherish. Could we hold up our heads as Americans if we allowed the time to come when she can no longer hold up hers?

Opportunities for Your Company.

OFFICIAL SPONSOR LIBERTY

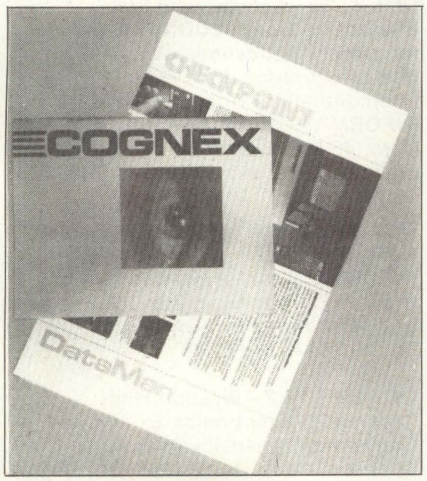
You are invited to learn more about the advantages of corporate sponsorship during the nationwide promotions surrounding the restoration project. Write on your letterhead to: The Statue of Liberty-Ellis Island Foundation, Inc., 101 Park Ave, N.Y., N.Y. 10178.



KEEP THE TORCH LIT

Save these monuments. Send your personal tax deductible donation to: P.O. Box 1986, New York, N.Y. 10018. **The Statue of Liberty-Ellis Island Foundation, Inc.**

Vision systems



Diverse vision system applications are covered in an eight-page color booklet. The Checkpoint intelligent inspection system gets special attention. **Cognex,** Needham, Mass. **Circle 410**

Logic analyzer

Full-color, 26-page data sheet features the HP 1630G logic analyzer, plus existent HP 1630A/D models. Application descriptions covering performance analysis, illustrations of screen displays, and a comprehensive data sheet are included. **Hewlett-Packard Co,** Palo Alto, Calif. **Circle 411**

Voltage regulators

Full-color brochure presents ac voltage regulator line; its eight pages feature a cost comparison chart determining energy savings, as well as specs and illustrations. **Powermark, Inc,** San Diego, Calif. **Circle 412**

Multibus product line

Catalog describes line of Multibus products that includes single-board computers, memory modules, I/O extenders, plus analog and peripheral controllers. Briefs on IEEE Standard 796 (Multibus) and IEEE P959 (iSBX) bus specs included. **Symbicon Assoc, Inc,** Amherst, NH. **Circle 413**

Heat sinks

Line of heat sinks and accessories is listed in an 80-page catalog. Straightforward guide to mounting methods is a catalog highlight. **AAALL, Inc,** Gilford, NH. **Circle 414**

Design and test instruments

Logic analyzers, digital multimeters, signal generators, and switching systems compose an updated instrument resource file. It features detailed product specs plus an overview of the test and measurement instrumentation spectrum. **Racal-Dana,** Irvine, Calif. **Circle 415**

Advanced I/O-intensive computing

Dimensions, a quarterly newsletter, covers the field of high speed I/O. Array processing, seismic data processing, image processing, computer graphics, data acquisition, and optical storage systems will be considered. **Aptec Computer Systems, Inc,** Portland, Ore. **Circle 416**

Public domain CP/M software

Over 250 public domain disk volumes from CP/M, SIG/M, and Pascal/Z user groups account for the 5000-plus files and programs listed in this CP/M public domain software catalog. Price is $5.00 in the U.S., California residents add tax. **Elliam Assoc,** 24000 Bessemer St, Woodland Hills, CA 91367.

Cables for IBM systems



Cables, cable assemblies, and interconnect accessories for IBM systems compose an updated reference guide. Divided according to specific systems and applications, PC XT and /36 cable varieties are spotlighted. **National Electric Cable,** Portland, Ore. **Circle 417**

From Z8 to Z80,000

A 40-page brochure describes line of microprocessors, development tools and associated peripherals. A preview is given of the forthcoming Z80,000 and Z800. **Zilog,** Campbell, Calif. **Circle 418**

Filters for PC boards

A 24-page catalog contains complete technical information on a line of feedthrough capacitors, ferrite bead inductors for noise suppression and spurious oscillation prevention, as well as emi suppression filters and noise filters. **Murata Erie North America, Inc,** Marietta, Ga. **Circle 419**

Half-height disk drives

A two-page data sheet offers details on the AST 96200 family of 61- and 103-Mbyte, 5¼-in. half-height rigid disk drives. **Advanced Storage Technology Inc,** San Jose, Calif. **Circle 420**

Telecommunication ICs

An expanded, eight-page pamphlet describes a telecommunication IC product line. Information on receivers/filters, modems, speech synthesizers and switched capacitor filter arrays merges with diagrams and tables containing maximum ratings and electrical characteristics. **Silicon Systems,** Tustin, Calif. **Circle 421**

Personal instrumentation

The MICROBASYS data acquisition system, which runs on the IBM PC, is detailed in a four-page pamphlet. Special treatment is given to communication utilities that allow simultaneous storage, presentation, and analysis of system data on a PC. **ADAC Corp,** Woburn, Mass. **Circle 422**

Sequential control

An application note dubbed "Economic Advantages of Sequential Control" describes the differences between traditional ladder-logic programming and direct sequential programming. **Control Technology Corp,** Westboro, Mass. **Circle 423**

Cross section cable ties

Heavy cross section cable ties for large size bundle diameters are profiled in this product bulletin. **Panduit Corp,** Tinley Park, Ill. **Circle 424**

CALENDAR



CONFERENCES

Invitational Computer Conf:
SEPT 6—Newton Marriott Hotel, Newton, Mass; **SEPT 25**—Southfield Sheraton Hotel, Southfield, Mich. **OCT 9**—Munich Sheraton Hotel, Munich, W Germany; **OCT 16**—Hilton International Hotel, Vienna, Austria; **OCT 23**—Hotel Executive, Milano, Italy. INFORMATION: Suzanne Hubner (U.S.) or Beatrice Labbe (Europe), B. J. Johnson & Assoc, Inc, 3151 Airway Ave, #C-2, Costa Mesa, CA 92626. Tel: 714/957-0171

SEPT 11-13—Midcon High Technology Electronics Exhibit & Convention, Dallas Convention Ctr, Dallas, Tex. INFORMATION: Kent Keller, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

SEPT 11-13—Mini/Micro Southwest Computer Conf & Exhibit, Dallas Convention Ctr, Dallas, Tex. INFORMATION: Kent Keller, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

SEPT 16-20—Compcon Fall, Hyatt Regency Crystal City, Arlington, Va. INFORMATION: IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

SEPT 19-21—Connector and Interconnection Technology Symposium, Disneyland Hotel, Anaheim, Calif. INFORMATION: Electronic Connector Study Group, Inc, PO Box 167, Fort Washington, PA 19034. Tel: 215/279-7084

OCT 1-4—AUTOFACT 6 Conf & Exhibit, Anaheim Convention Ctr, Anaheim, Calif. INFORMATION: Society of Manufacturing Engineers, PO Box 930, Dearborn, MI 48121. Tel: 313/271-0777

OCT 2-4—Northcon High Technology Electronics Exhibit & Convention, Seattle Center Coliseum, Seattle, Wash. INFORMATION: Nancy Hogan, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

OCT 2-4—Mini/Micro Northwest Computer Conf & Exhibit, Seattle Center Flag Pavilion, Seattle, Wash. INFORMATION: Nancy Hogan, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

OCT 8-10—9th Conf on Local Area Networks, Minneapolis, Minn. INFORMATION: Harvey A. Freeman, Architectural Technology Corp, PO Box 24344, Minneapolis, MN 55424.

OCT 8-10—ACM Annual Conf: The Fifth-Generation Challenge, San Francisco Hilton Hotel, San Francisco, Calif. INFORMATION: Karen Duncan, Chairman, ACM '84, 15 Parsons Way, Los Altos, CA 94022.

OCT 10-12—Design Automation Workshop, East Lansing, Mich. INFORMATION: Harry Hayman, IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 10-12—LOCALNET '84, Sheraton Harbor Island Hotel, San Diego, Calif. INFORMATION: Online Conferences, Inc, Suite 1190, 2 Penn Plaza, New York, NY 10121. Tel: 212/279-8890

OCT 11-12—Peripheral Array Processor Conf, Copley Plaza Hotel, Boston, Mass. INFORMATION: Charles A. Pratt, Simulation Councils, Inc, PO Box 2228, La Jolla, CA 92038. Tel: 619/459-3888

OCT 15-17—The Future of Optical Memories, Loew's Summit Hotel, New York, NY. INFORMATION: Joanna Spilman, Technology Opportunity Conference, PO Box 14817, San Francisco, CA 94114. Tel: 415/626-1133

OCT 15-18—Conf on Ada Applications & Environments, Sheraton Midway Hotel, St. Paul, Minn. INFORMATION: Harry Hayman, IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 16-18—Unix Operating System Expo & Conf, Sheraton Centre Hotel, New York, NY. INFORMATION: Robert P. Birkfield, National Expositions Co, Inc, 14 W 40th St, New York, NY 10018. Tel: 212/391-9111

OCT 16-19—Int'l Test Conf (Cherry Hill '84), Franklin Plaza Hotel, Philadelphia, Pa. INFORMATION: Harry Hayman, IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 22-24—13th Annual Workshop on Applied Imagery Pattern Recognition, Sheraton Inn-Washington Northwest, Silver Spring, Md. INFORMATION: IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 22-25—ISA '84, Astrohall, Houston, Tex. INFORMATION: Instrument Society of America, 67 Alexander Dr, Research Triangle Park, NC 27709. Tel: 919/549-8411

OCT 22-26—Annual Industrial Electronics Conf (IECON '84), on Industrial Applications of Microelectronics, Keio Plaza Intercontinental Hotel, Tokyo, Japan. INFORMATION: Frank A. Jur, Bechtel Corp, 45 Fremont St, MS 45/17A26, San Francisco, CA 94109. Tel: 415/768-3023

OCT 30-31—Flat Information Display Conf, Red Lion Inn, San Jose, Calif. INFORMATION: Murray Disman, Int'l Planning Information, Inc, 164 Pecora Way, Portola Valley, CA 94025. Tel: 415/854-7306

OCT 30-NOV 2—Wescon High Technology Electronics Exhibit and Convention, Anaheim Convention Ctr, Anaheim, Calif. INFORMATION: Dale Litherland, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

NOV 12-15—IEEE Int'l Conf on Computer Aided Design, Santa Clara, Calif. INFORMATION: John A. Domiter, Bell Telephone Labs, 4K523, Holmdel, NJ 07733. Tel: 201/949-6675

NOV 13-17—Elektronica, Munich Trade Fair Center, Munich, W Germany, INFORMATION: Kallman Assoc, 5 Maple Ct, Ridgewood, NJ 07450. Tel: 201/652-7070

NOV 16-17—Forth Interest Group Convention, Hyatt Palo Alto, Palo Alto, Calif. INFORMATION: Forth Interest Group, PO Box 1105, San Carlos, CA 94070. Tel: 415/962-8653

DEC 3-9—Int'l Microcomputer Conf & Display 1984, China, The Guangdong Scientific Hall, Guangdong, PRC. INFORMATION: Meridian Technology Exhibitions Ltd, Rm 1201 Kai Tak Commercial Bldg, 317 Des Voeux Rd, C, Hong Kong

DEC 11-13—Fifth-Generation and Supercomputer Symposium, Rotterdam, The Netherlands. INFORMATION: Rotterdam Tourist Office, Stadhuisplein 19, 3012 AR Rotterdam. Tel: (010) 14 14 00

SHORT COURSE

NOV 19-20—Applications-Oriented Approach to Artificial Intelligence, George Washington Univ, Washington, DC. INFORMATION: Continuing Engineering Education Program, George Washington Univ, Washington, DC 20052. Tel: 202/676-6106

Why this magazine and more than 1,000 others let us go over their books once a year.

Some magazines, we're sorry to say, keep their readers undercover. They steadfastly refuse to let BPA (Business Publications Audit of Circulation, Inc.) or any other independent, not-for-profit organization audit their circulation records.

On the other hand, over 1,000 publications (like this one) belong to BPA. Once a year, BPA auditors examine and verify the accuracy of our circulation records.

This audit provides the name, company, industry and job title of every reader each publication reaches. The information helps advertisers to determine if they are saying the right thing to the right people in the right place.

It also helps somebody else important: you. Because the more a publication and its advertisers know about you, the better they can provide you with articles and advertisements that meet your informational needs.

BPA. For readers it stands for meaningful information. For advertisers it stands for meaningful readers. Business Publications Audit of Circulation, Inc.
360 Park Ave. So., New York, NY 10010.

BPA
MEDIA INTELLIGENCE

DO YOU OFTEN WISH THAT THE SMITH COMPANY WAS THE SMITH & SMITH & SMITH & SMITH COMPANY?

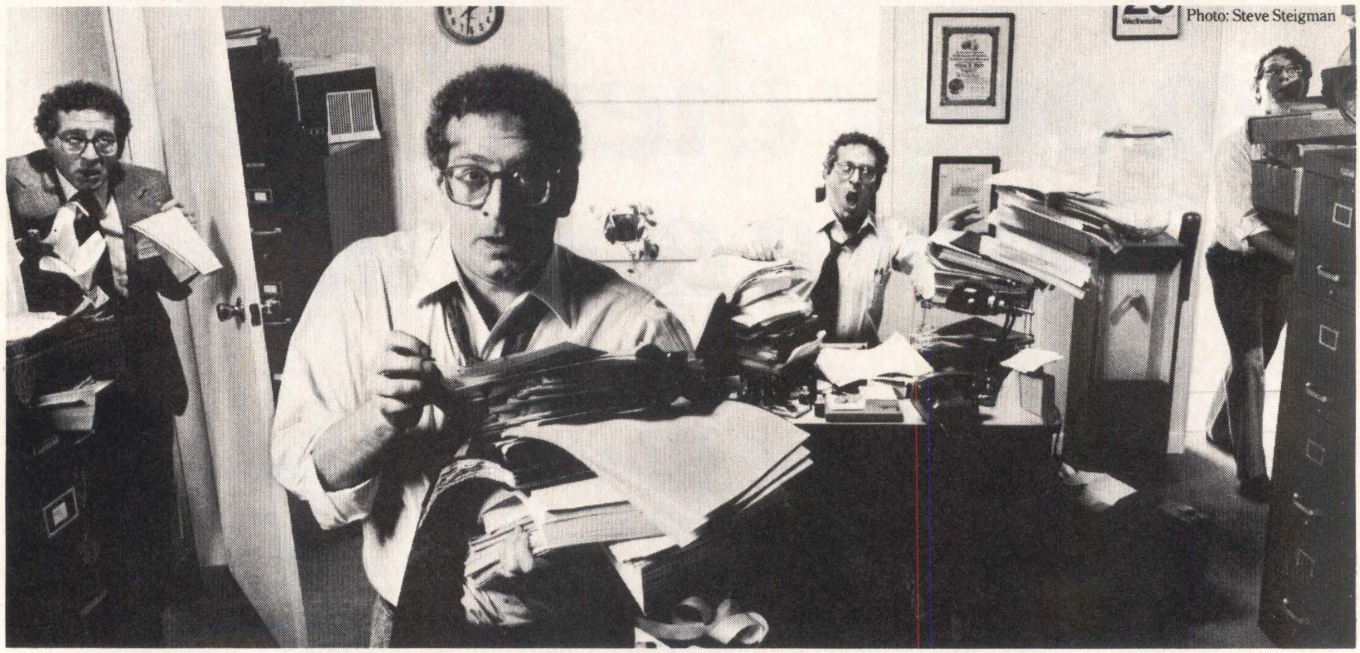


Photo: Steve Steigman

Business people often find themselves doing so many things that just one of them isn't enough.

It happens all the time. Because you start out in a one-person shop and soon you need three to handle the job.

If you're in that sort of a predicament, we at the National Alliance of Business would like to offer a suggestion, and a solution.

Hire some of the needy and disadvantaged young people of America this summer. And let them put their muscles and minds to work. For you.

Giving them work will not only relieve some of the pressure on you now, but it will make it easier for them to become part of the full-time labor force later on.

Hiring them can also help you in another very important way. Because a business that hires economically disadvantaged youth during the summer may get as much as an 85% tax credit on the first $3,000 of wages you pay them.

Write the National Alliance of Business, P.O. Box 7207, Washington D.C. 20044 for more information. And support your local summer-jobs-for-youth programs.

You'll be doing something for yourself. You'll be doing something for your business.

And you'll be doing something for your community and the needy youth of America, too.

LET'S GET ALL OF AMERICA WORKING AGAIN.

A PUBLIC SERVICE
OF THIS PUBLICATION &
THE ADVERTISING COUNCIL. Ad Council



DESIGNER'S BOOKCASE



SOFTWARE TESTING TECHNIQUES

By *Boris Beizer*

A storehouse of sophistication methods for test case design and execution. The author clarifies graph theory and presents alternate program models that ensure flexibility in testing processes. He shows how to utilize McCabe's metric, decision tables and boolean algebra, and much more.

$29.95 **Circle 445**

COMPUTER METHODS FOR CIRCUIT ANALYSIS AND DESIGN

By *Jiri Vlach and Kishore Singhal*

A storehouse of computer-aided techniques for analyzing and designing a wide range of linear and nonlinear circuits, including analog networks, analog filters, and switched-capacitor networks. Provides tips on how to write programs for new applications as well as how to modify existing programs.

$42.50 **Circle 446**

SYNTHESIS OF DIGITAL DESIGNS FROM RECURSION EQUATIONS

By *Steven D. Johnson*

Defines a circuit description language that uses systems equations to state connectivity—systems that are, in fact, applicative programs that compute the logical behavior of the circuit described. An interpreter is then presented through which both specifications and target descriptions can be executed, allowing engineers to experiment directly with the design notation without having to translate into a simulation language or construct a physical prototype.

$30.00 **Circle 447**

HOW TO ORDER:

15-DAY FREE EXAMINATION

(U.S. AND CANADA ONLY)

Simply circle the appropriate number(s) on the Reader Inquiry Card at the back of this magazine.

Your book will be sent to you for your **15-day free trial.** If you are satisfied, keep the book and an invoice will follow. Otherwise return the book by the end of the 15-day period, and owe nothing.

MICROPROCESSORS IN INDUSTRY

By *Michael Hordeski*

Here is the first book to integrate all aspects of microprocessor systems design and control theory for industrial applications. In view of recent developments in distributed processing and computer networks, readers will especially appreciate chapter on the latest 16-bit microprocessors, bubble memory systems, software design techniques, transmission techniques and communication protocols.

$49.50 **Circle 448**

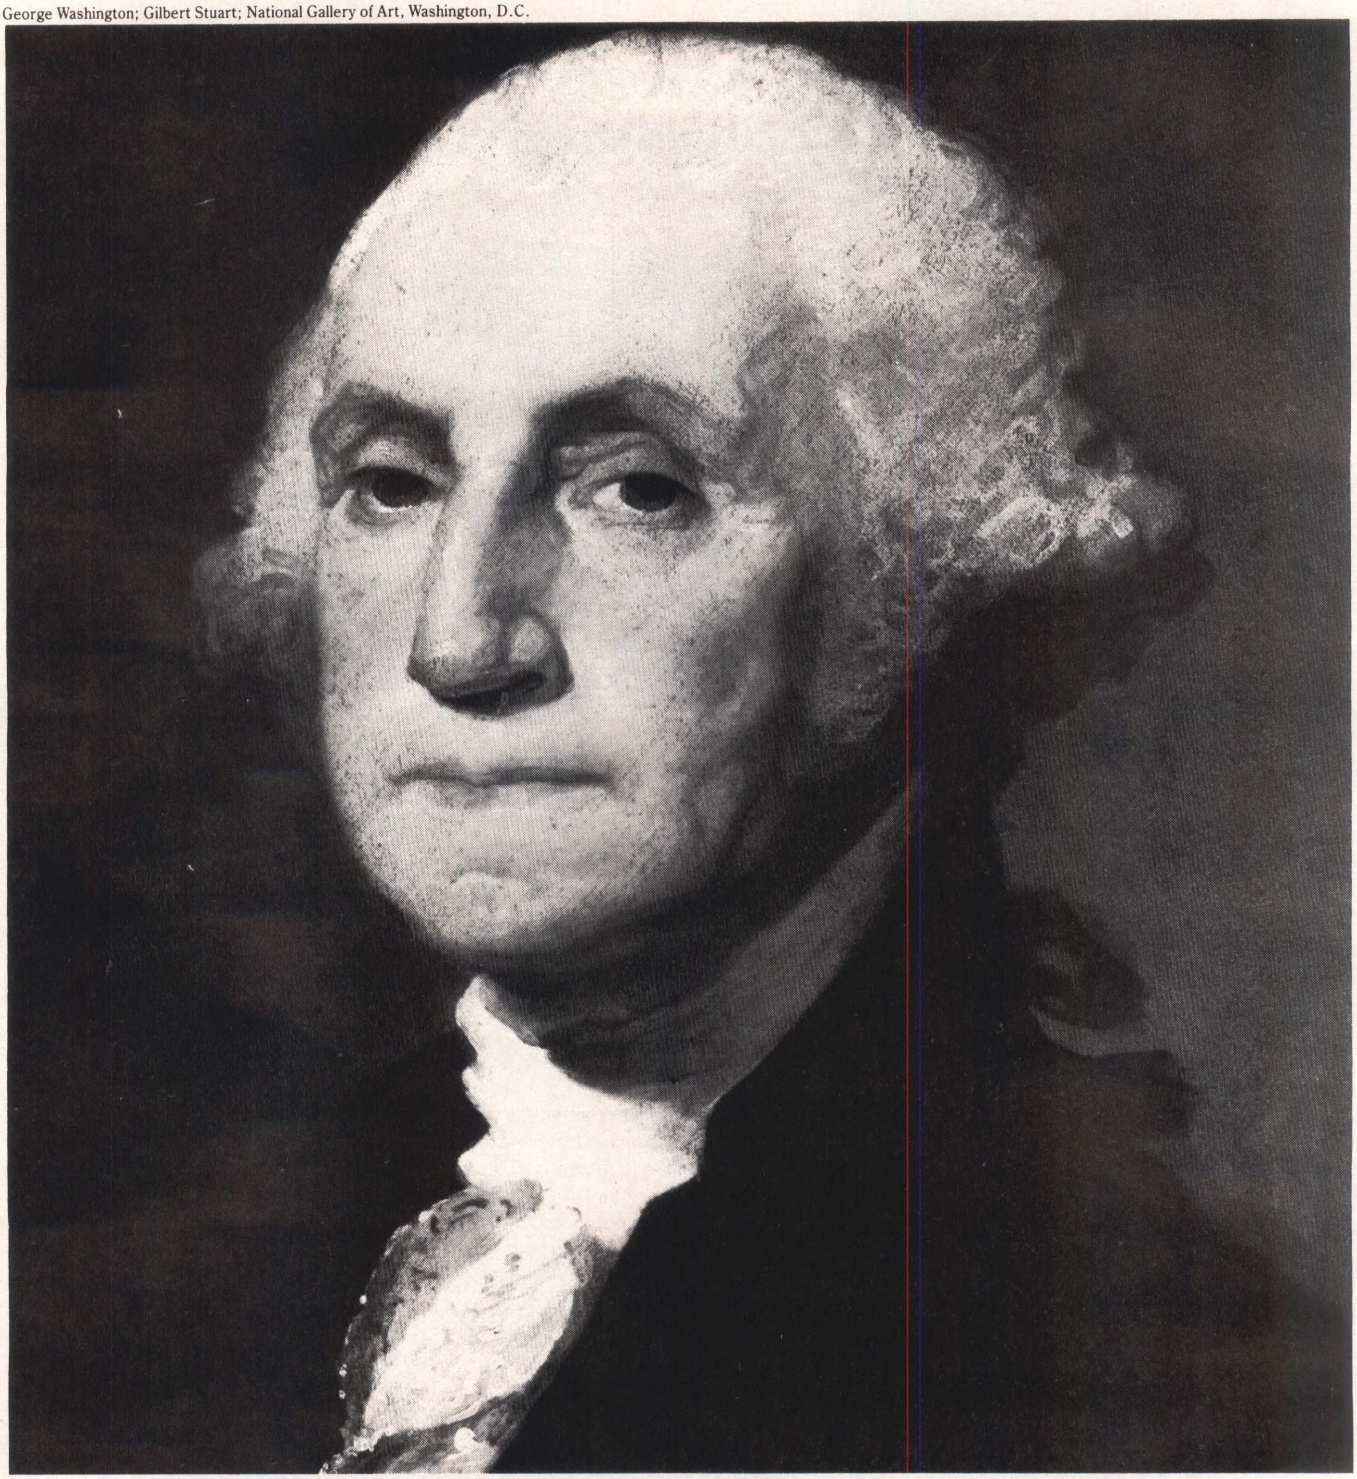
LOGICAL DATA BASE DESIGN

By *Robert M. Curtice and Paul E. Jones*

This far-reaching guide provides a workable framework for appraising a set of concepts about data and data base as well as special time-saving tips on how to document a logical design. Examples promote easy application of information to on-the-job situations.

$34.50 **Circle 449**

George Washington; Gilbert Stuart; National Gallery of Art, Washington, D.C.



Would the winter he spent at Valley Forge have cost him a promotion in your company?

For over three centuries citizen soldiers have left their homes and jobs to defend America as George Washington once did. They still do today. They're the members of the Guard and Reserve. Their readiness is vital to our defense.

Be fair to them when their time comes to go on duty. Don't penalize them when promotions and raises are due. Our country's future still depends upon our citizen soldiers. They depend upon you, their employers, now.

Find out how you can support today's Guard and Reserve. Write Employer Support, Arlington, VA 22209. Or call 800-336-4590.

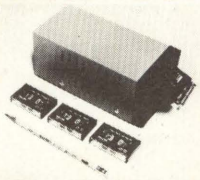
Protect their future while they protect yours.

ESGR EMPLOYER SUPPORT OF THE GUARD & RESERVE Ad Council A Public Service of This Publication

SYSTEM SHOWCASE

CIRCLE 475 for rates and information, OR CALL: SHIRLEY LESSARD 800-225-0556, in MA 617-486-9501

RS232 or 8 BIT PARALLEL MINI CASSETTE SYSTEM



- LOW COST – $388 @ 100 UNITS*
- Microprocessor Controlled
- >200K Bytes Formatted/Tape
- Variable Baud Rate (110-9600)
- Error Rate 1 in 10 to the 9th
- Small — 3" x 3.5" x 5" W/O Case*

21 OTHER CARTRIDGE, CASSETTE, FLOPPY SYSTEMS WITH STORAGE TO 12 MEGABYTES

Basic Drive with R/W Electronics Start at $150

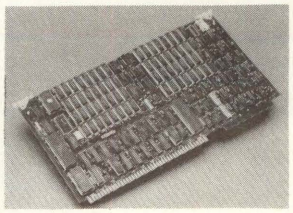
ANALOG & DIGITAL PERIPHERALS INC

ADPI

815 Diana Drive
Troy, Ohio 45373
513/339-2241
TWX 810/450-2685

Branch Off: Oklahoma City OK – Factory: Yucca Valley,CA

CIRCLE 476

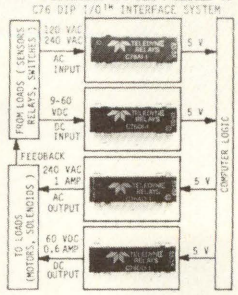


DISK CACHE FOR INTEL DEVELOPMENT COMPUTERS

Improves performance 3 times without any changes. For Series 2 or 3, MDS 800 and NDS II work stations. Works with single or double density floppies or hard disk. Automatic allocation of 480 K-Bytes of cache memory. **Built in 196 K-Byte print spooler.** The model OR-88C is $2850 each, and in stock. **ORIGIN, INC,** 9136 Gibson St, LA, CA 90034. Tel: (213) 202-0772. TLX 704809.

CIRCLE 478

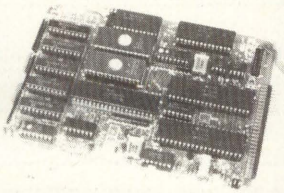
TELEDYNE RELAYS



C76 SERIES SOLID STATE DIP I/O™

Modules are designed for use in computerized control systems for noise-free isolated interfacing of computer logic elements to harsh industrial environments. Modules employ custom microcircuits in a TO-116 DIP. Low "Per-point" cost. Meets VDE spacing and voltage requirements. Features CMOS and TTL compatibility, optical isolation, high noise immunity and an ENABLE function on input modules. $8.70 to $9.65 (depending on function) at production quantities. **TELEDYNE SOLID STATE PRODUCTS,** 12525 Daphne Ave, Hawthorne, CA 90250. Tel: (203) 777-0077.

CIRCLE 479

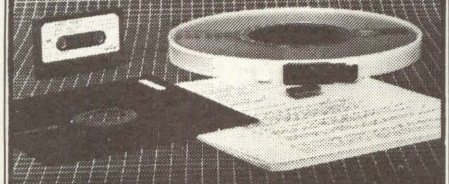


10809 STD BUS SINGLE CARD COMPUTER/CONTROLLER

uses 6809 (E)MPU. Provides cost effective multicard capability-2K CMOS RAM may be battery backed (power fail/interrupt on board), sockets for 28 pin EPROM/EEPROM, RS232/423 serial I/O, 32 lines parallel I/O, 4 timers. Other 10000 SERIES MICRO-COMPUTER MODULES include 10812 6502 SCC/C, 10764 Universal Non-Volatile Memory, 10611 Universal TTL I/O (w/8 ports-64 lines), 10301 Versatile Video Display Card. Custom design/mfg. to your requirements. Call today. **ENTERPRISE SYSTEMS CORP,** Box 698 Dover, NH 03820. Phone (603)-742-7363.

CIRCLE 480

6800 Family Software

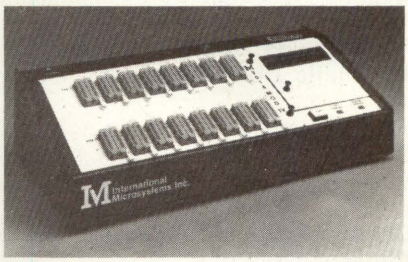


Software for 6800/01/05/09 microprocessors.
SOFTWARE: C compiler, Pascal, multitasking DOS, macro-assemblers, FORTRAN, 16K BASIC
FIRMWARE: Fantom monitor/debugger, math library, 4K BASIC
CROSS-SOFTWARE: Assemblers, PL/W Compilers, Linker, Simulators

WINTEK

Wintek Corp.
1801 South Street
Lafayette, IN 47904
317-742-8428

CIRCLE 481



MULTI-SET EPROM PROGRAMMER

IM3016 programs up to 16 DIFFERENT IMAGES in one cycle. Ideal for engineering and/or production environments using large EPROM SETS. Electrically isolated sockets, one-key control and APPROVED fast algorithms. Software selection for all 5V EPROMs and CMOS versions including 2764, 2764A, 27128, 27128A & 27256. **INTERNATIONAL MICROSYSTEMS, INC.** 11554 C Avenue, Auburn, CA 95603. Tel: (916) 885-7262.

CIRCLE 482

LOW COST ADVERTISING

You can market your products to over 90,000 systems builders

- Hardware
- Software
- Services
- Consulting

For only $655.00 your ad will be typeset, laid out, and appear in this space. Sales leads will automatically be sent to you as we receive them.

Start today! Call Shirley Lessard for details on how to start a low cost advertising program. Your sales results will tell you that it is the best decision you have ever made.

COMPUTER DESIGN
(800)225-0556 (Outisde Mass)
(617)486-9501

CIRCLE 483

8048

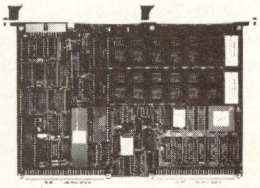
- SIM48* software simulates and debugs 8048 programs
- Runs on IBM-PC*, CP/M-80*, MS-DOS*
- Internal timer, interrupt, inst set, external RAM, 4K ROM supported
- One year FREE updates
- $250., manual alone $20., shipping and handling included
- Formats — PC DOS DSDD, CP/M-80 8" SSSD, Kaypro II SSDD

Logical Systems

6184 TEALL STATION
SYRACUSE, NY 13217
(315)457-9416

CIRCLE 484

OMNIBYTE™



OB68K/VMEI SINGLE BOARD COMPUTER ON THE VME BUS

- 12.5MHz 68000R12 16/32 bit CPU • (8) pairs of 28-pin sockets for RAM or ROM • (2) RS-232C serial ports using (1) 68681 DUART • (2) 8-bit parallel I/O ports using (1) 68230 PI/T • System controller functions are supported • (7) Prioritized bus or auto vectored prioritized interrupts
- Two year limited warranty.

OMNIBYTE CORPORATION
245 W. Roosevelt Road
West Chicago, IL 60185
(312) 231-6880

CIRCLE 485

AD INDEX

Able Computer....183
ACC....C4
Advanced Micro Devices....72, 73
AMCC....115
American Hoechst Plastics....179
AMF-Potter & Brumfield....40, 41
AMP....152, 153
Ampex....70, 71
Amphenol Products, an Allied Co....36, 37
Analog & Digital Peripherals....279
Atlantic Research Corp....66
AT&T Technologies....142, 143, 151
Audiotronics....135
Automatic Coil Corp....48

Ball Electronic Systems....34
Burr-Brown Corp....154, 204
Business Computing....257, 258, 259

Cherry Electrical Products....22
Chrislin Industries....16
Comchek International....141
Computer Products....256
Contact Electronics....61
CSPI....28
Cubit....242

Dataram Corp....7, C3
Davidge Corp....67
Deltron....268
Dialight Corp....47
Digi-Data Corp....164
Digital Engineering....230, 231
Diversified Technology....85

Eagle Magnetic Co....173
Elpac Power Systems....81
Emulogic....57, 149
Engineering Automation Systems....188
Ensign-Bickford Optics Co....46
Enterprise Systems Corp....279
Excelan....9

Fairchild Digital....35
Federal Reserve Bank of Dallas....270
Fujitsu Microelectronics....113

General Electric Plastics....49, 238, 239
Gould....88, 89, 200, 201
Graphic Strategies....56

Hewlett Packard....117, 171
Hitachi....232
Hitachi America Ltd....216, 217
Houston Instrument....C4
Hunter & Ready....235

IBM....39
IBIS Systems....127
IMC/Hansen....246
Inmos Corp....202, 203
Intel....20, 21, 62, 63, 158, 159
Interface Technology....97
International Microsystems....279
Interphase....55
Invitational Computer Conference....240, 267
Iomega Corp....132
ITT Schadow....193

Keltron Corp/High Voltage....16
Kennedy Co....1
Konan Corp....109

Lear Siegler....125
Logical Systems Corp....279

Marinco Computer Products....174, 175
Megatek Corp....17, 93
Memory Protection Devices....245
Mentor Graphics....86, 87
Microbar....243
Micro Switch....220, 221
Mini/Micro....116
Molex....282
Mostek....42, 43, 191
Motorola Semiconductor Products....25, 26, 27
Multi Solutions....12
Multiwire Corp....181
Microtek Lab....68

NCR Corp....162, 186, 187
NEC Electronics....C2

Oaysis....45
Okidata....83
Omnibyte Corp....279
Origin....279

Panasonic....195, 215
Pertec Peripherals Corp....222
*Plessey Microsystems....260
Princeton Graphics....176
PrintaColor Corp....249
Pyramid Technology....120

Quantum....58, 59

Rockwell International....104
Rogers Corp....69
RPM Corp....173
RTCS....263

Samtec....169

Scientific Micro Systems....229
Seiko Instruments....207
Shugart Associates....14, 15, 74, 75
Signetics....253, 254, 255
Silvar-Lisco....90
Solid State Scientific....111
Southern Computer Corp....24
Stackpole Components Co....160

Tandon Corp....10, 11
Tauber Electronics....271
Tektronix....33
Teledyne Solid State Products....279
T&B Ansley Electronics....102
TII Electronics....144
Toshiba America....98, 99
TRW....76
TRW/LSI Products....199

Universal Data Systems....64

VCA Corp....266
Vermont Research Corp....250
Vesatec, a Xerox Co....50, 51
Versitron....211
VLSI Technology....256a-256d

Western Digital Corp....218, 219
Wintek Corp....279
Wyse Technology....52, 225

Xebec....130, 131
Xicor....100

ZAX Corp....118, 119
Zilog....19

*International Issues Only

The AD INDEX is published as a service. The publisher does not assume any liability for errors or omission.

CD* SALES OFFICES

Home Office 119 Russell St., Littleton, MA 01460 (617) 486-9501

Marketing Director, Robert Billhimer
Ad Traffic Coordinator, Debra Friberg
Systems Showcase, Shirley Lessard
Postcard Deck, Shirley Lessard
Classified/Recruitment, Shirley Lessard
List Rental, Robert Dromgoole

New England/Upstate New York
Barbara Arnold
119 Russell St.
Littleton, MA 01460
(617) 486-9501

Mid-Atlantic/Southeast
Richard V. Busch
40 Stony Brook Lane
Princeton, N.J. 08540
(609) 921-7763
Eleanor Angone
74 Brookline Ave.
E. Atlantic Beach, NY 11561
(516) 432-1955

Midwest & Colorado
Berry Conner
88 West Schiller St., Suite 2208
Chicago, IL 60610
(312) 266-0008

Southwest
Steve Lassiter
1200 S. Post Oak Blvd.
Houston, TX 77056
(713) 621-9720

Southern California
Buckley/Boris Associates
Tom Boris, John Sabo
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
(714) 957-2552

Northern California
Buckley/Boris Associates
Tom Boris, John Sly
920 Yorkshire Drive
Los Altos, CA 94022
(415) 964-4232

Northwest
Buckley/Boris Associates
Tom Boris
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
(714) 957-2552

International
International Sales Manager
Eric Jeter
1200 S. Post Oak Blvd.
Houston, TX 77056
(713) 621-9720

U.K. and Scandinavia
David Betham-Rogers, David M. Levitt
6th Floor, Alliance House
12 Caxton Street
Westminster, London SW1H OQS
Tel: 01-222 0744
Telex: 919775

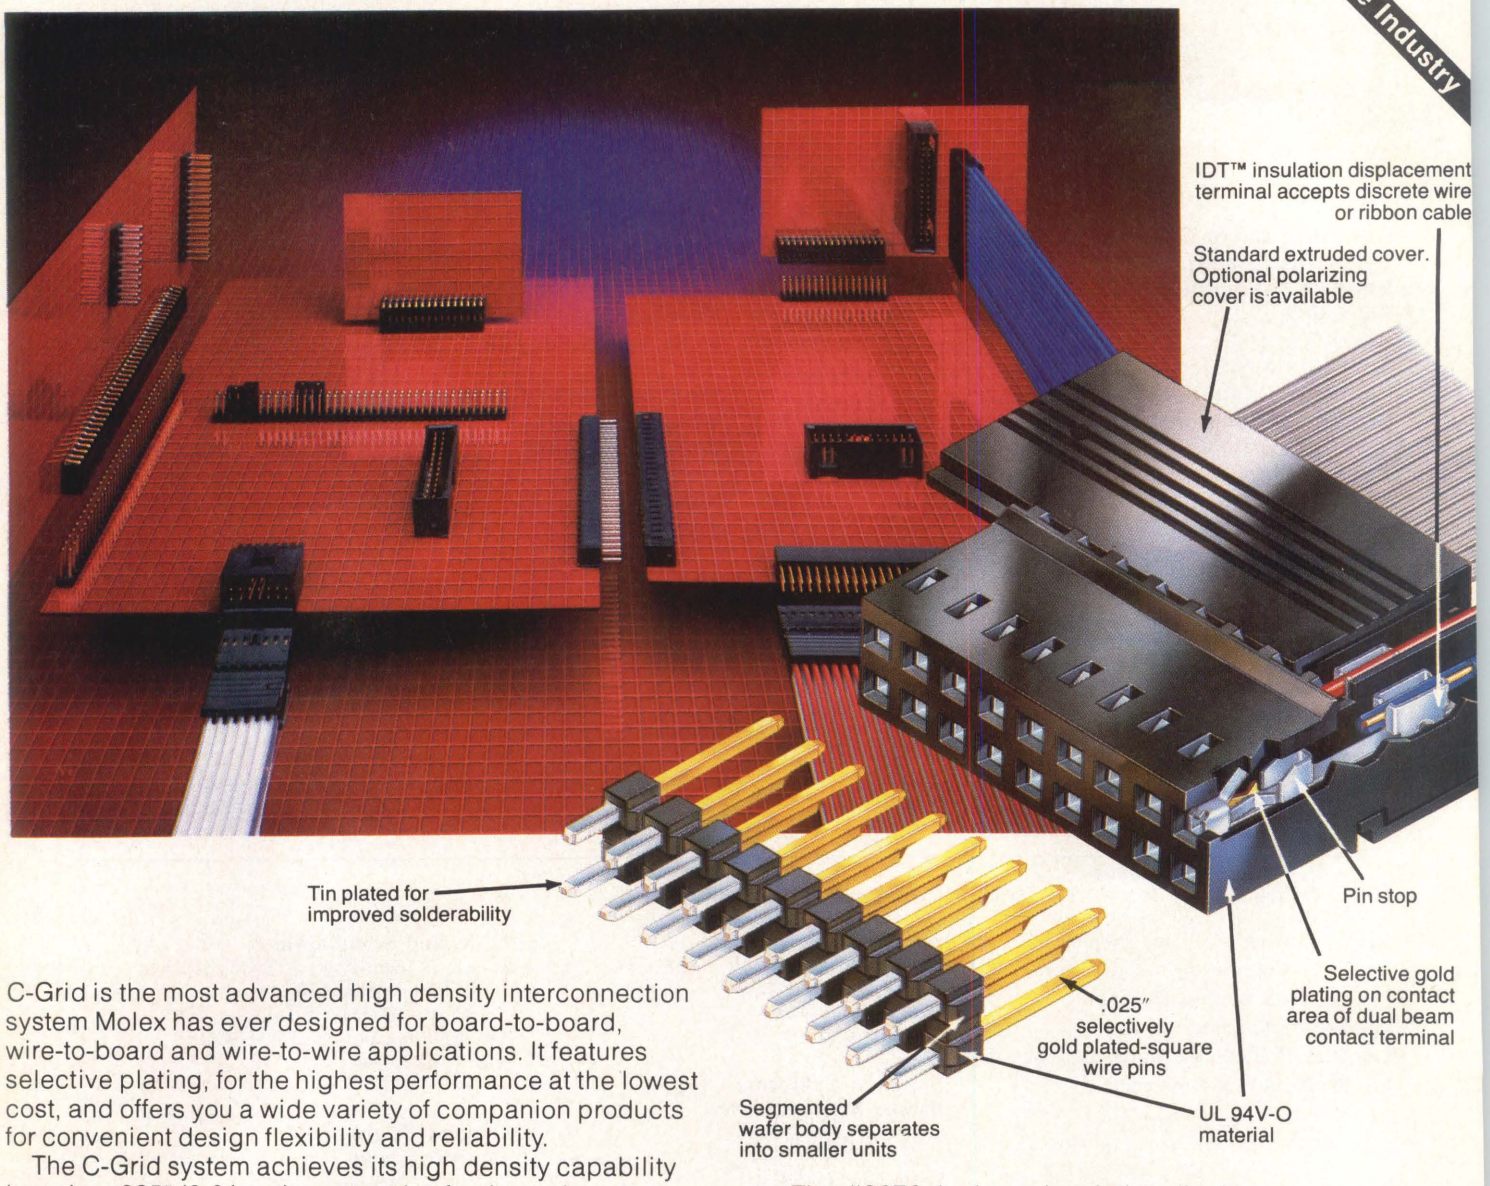
France, Belgium and S. Switzerland
Daniel R. Bernard
247, Rue Saint Jacques
75005 Paris
Tel: (1) 354.55.35
Telex: 250 303

Holland, Austria, W. Germany, Switzerland & Eastern Europe
Heinz Gorgens
Parkstrasse 8a
D-4054 Nettetal 1—Hinsbeck (F.R.G.)
Tel: (0 21 53) 8 99 88/89
Telex: (17) 2153310

Japan
Sumio Oka & Shigeo Aoki
International Media
Representatives Ltd.
2-29 Toranomon 1-chome
Minato-ku, Tokyo 1Q5 Japan
Tel: 502-0656
Telex: J22633 *(TM)

C-Grid: the Molex .100" x .100" PCB interconnection system.

Shortest Leadtimes in the Industry



C-Grid is the most advanced high density interconnection system Molex has ever designed for board-to-board, wire-to-board and wire-to-wire applications. It features selective plating, for the highest performance at the lowest cost, and offers you a wide variety of companion products for convenient design flexibility and reliability.

The C-Grid system achieves its high density capability by using .025" (0,64mm) square pins for the male connector parts, set in a .100" x .100" (2,54 x 2,54mm) matrix. The pins can either be placed directly in the board or used in volume with wafer bodies.

To load loose pins into the PC board, Molex offers a patented single- or Multi-Pinsetter® capable of up to 156,000 insertions per hour—the fastest on the market.

Shrouded and unshrouded headers are available in straight and right-angle wafer bodies.

The #8676 dual-row, insulation displacement connectors are stackable, end-to-end, and feature IDT™ terminals with mass termination application tooling available to help you achieve greater cost savings in your assembly operations.

Molex closed-or open-end shunts make it easier and less expensive to retrofit board circuitry by avoiding DIP switches.

Molex also offers a female connector (7990-series) to be soldered to one board and mated with the .025" (0,64mm) pins on another. Its features include low mating force, ease of insertion and improved solderability.

For more information on our C-Grid interconnection system, contact the Molex office nearest you.

First in Customer Service ...Worldwide

molex

CIRCLE 134

Corporate Headquarters: 2222 Wellington Court, Lisle, Illinois 60532 Phone: (312) 969-4550 Telex: 27-0072/25-4069
European Headquarters: Molex House, Church Lane East, Aldershot, Hants, England GU11 3ST Phone: (0252) 318221 Telex: 851858988
Far Eastern Headquarters: 5-4, 1-Chome Fukami-Higashi, Yamato-Shi, Kanagawa Pref., Japan Phone: (462) 614500 Telex: 781-03872486

...from DATARAM



MOTOROLA VERSAbus 1.0 MB ADD-IN

Not just more capacity; more capability

Motorola's 68000 is a winner, and using this popular microprocessor in a VERSAbus/VERSAboard configuration is a smart move. Dataram's single-board 1.0 MB DR-680 can make it even smarter. Increased single-board memory capacity means lower power, less space, higher reliability, and lower cost...and the DR-680 provides a lot more!

More speed: Ability to perform match cycles reduces access/cycle times by allowing immediate reading of data registers when adjacent words/bytes are accessed.

Advanced error handling: The DR-680 provides an on-board control and status register (CSR) which allows program control of ECC functions and contains the diagnostic information required for error analysis. The CSR can be read or written via the VERSAbus. Additionally, the DR-680 greatly increases reliability by performing error "sniffing" and error "scrubbing" during refresh operations.

Compatibility: Dataram's DR-680 is compatible with Motorola's EXORmacs, IBM System 9000, and all other VERSAbus-based systems.

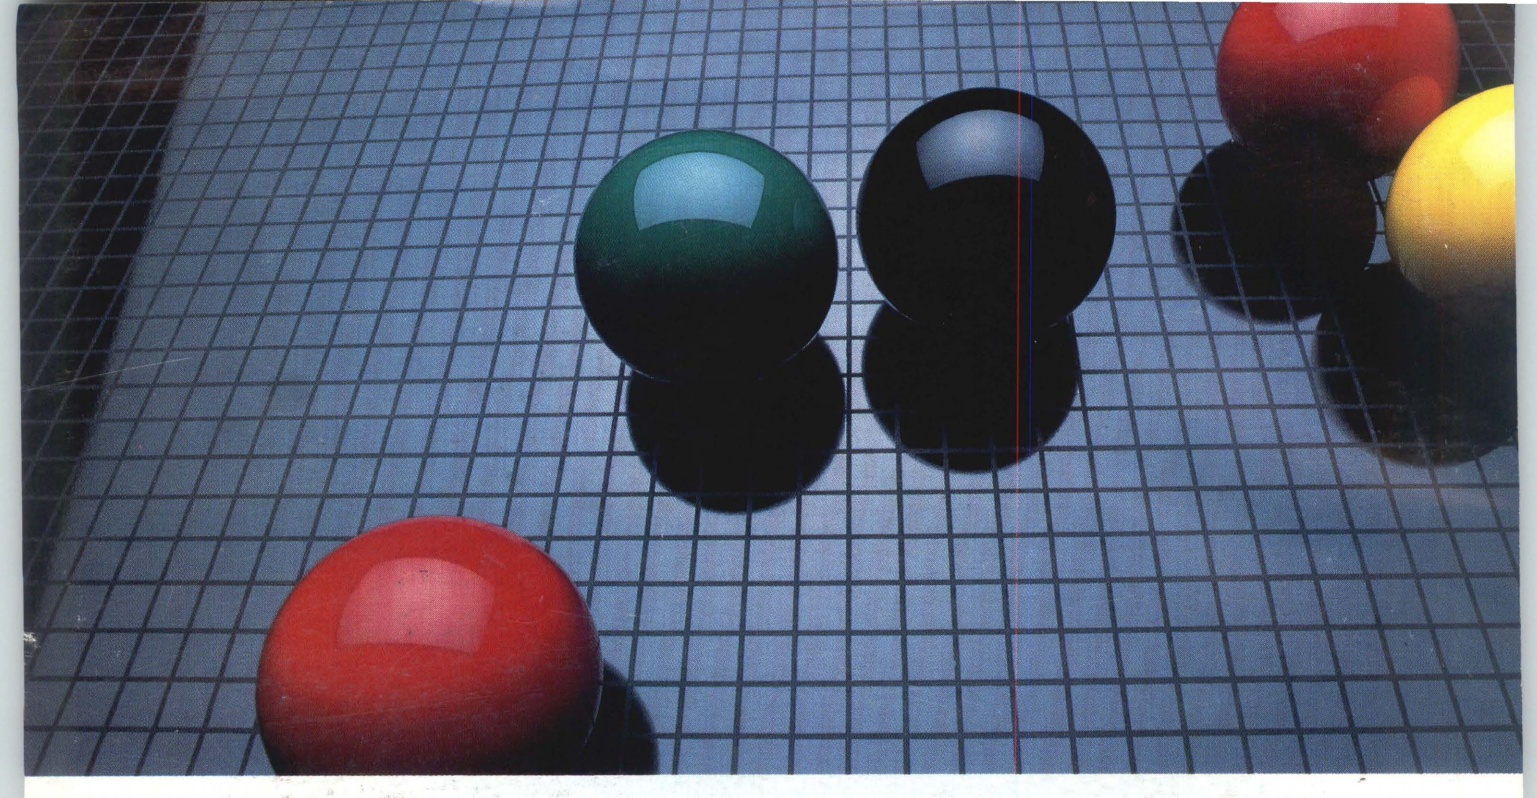
EXORmacs, VERSAboard and VERSAbus are registered trademarks of Motorola, Inc.

For more information, call Dataram today at (609) 799-0071.

DATARAM

Dataram Corporation ☐ Princeton Road ☐ Cranbury, NJ 08512 ☐ (609) 799-0071 ☐ TWX: 510-685-2542

CIRCLE 135



CREATE A UNIFIED, MULTI-VENDOR NETWORK.

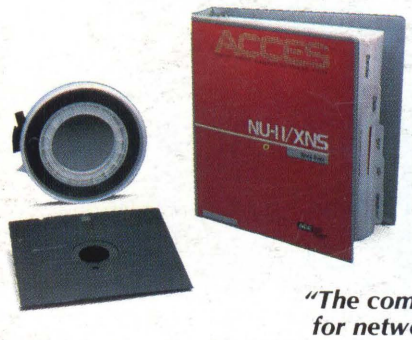
Combine the strengths of multiple computer vendors through the capabilities of the ACC Exchange System (ACCES).

All layers of the OSI Reference Model are provided using ACCES network protocols and application layer packages. Using industry standards, ACCES provides the structure and the capability for true distributed processing.

File transfer between different operating systems, file management, file access to the record level, and uniform terminal access — across the entire network.

Talk to ACC and learn about the ACCES Product Line. To receive more information on ACCES and on ACC's network solutions, give us a call today (805) 963-9431.

Advanced Computer Communications.
720 Santa Barbara Street
Santa Barbara, CA 93101
TWX 910 334-4907



ACC

Advanced Computer Communications

"The company relied on most often for network solutions"

CIRCLE 136

© 1984 Advanced Computer Communications